

09/270733  
03/16/99

Inventor(s): David MacDonald DELANEY et

## VIRTUAL PRIVATE NETWORKS AND METHODS FOR THEIR OPERATION

- RJH11/97

## VIRTUAL PRIVATE NETWORKS & METHODS FOR THEIR OPERATION

### Field of Invention

This invention relates to Virtual Private  
5 Networks (VPNs) and to methods for their operation. More particularly this invention relates to methods and apparatus that enable Network Service Providers (NSPs) to provide virtual private LAN interconnect services to large groups of customers.

10

### Background of Invention

Most large businesses operate LANs at several sites to meet their data communications needs. The businesses lease dedicated circuits from NSPs to connect  
15 their LANs into Wide Area Networks (WANs). Because distinct customers of the NSP lease distinct dedicated circuits, their WANs are isolated from another, thereby meeting data security requirements.

The dedicated circuits are available in fixed  
20 bandwidths (e.g. DS1, DS3). Customers must lease a dedicated circuit that meets their maximum bandwidth requirements. Because typical data traffic is bursty, whereas the dedicated circuits provide a fixed bandwidth at all times, the dedicated circuits are frequently  
25 operating below capacity. Consequently, customers typically pay for more dedicated circuit capacity than they would need if the NSP's network capacity could be shared more efficiently among customers while preserving the required isolation between networks of distinct  
30 customers.

The IEEE 802.1 standard defines a protocol that enables an Ethernet LAN to be partitioned into multiple Virtual LANs (VLANs), each VLAN being isolated from the other VLANs. Large businesses typically use the IEEE  
35 802.1 protocol to partition their LANs into VLANs for distinct interest groups within the business.

The IEEE 802.1 standard requires that a header of each frame of data carry a VLAN tag that identifies the VLAN for which the data frame is intended. Switches (or "bridges") of the LAN read the header and route the data frames to only those ports which, according to routing tables (or "filter databases") stored at the switches, are participating in that VLAN. The 12 bit capacity of the VLAN tag specified by the IEEE 802.1 standard limits the number of distinct VLANs to 4095. NSPs need to support many more than 4095 distinct customers on a shared network.

#### Summary of Invention

In this specification, the terms "switch", "switching element", "router" and "routing device" are intended to include any device providing switching or routing functionality including, but not limited to, switches and routers.

This invention seeks to provide methods and apparatus that enable a NSP to provide a very large number of VLANs on shared network facilities.

Embodiments of the invention may use extensions to Ethernet protocols so that existing Ethernet technology and familiarity with Ethernet in the data communications industry can be leveraged to provide VLAN capability for a large number of customers at low acquisition cost and low operating cost.

One aspect of the invention provides a method of routing packets through a communications network having a plurality of distinct sets of virtual ports. No virtual port belongs to more than one of the distinct sets. In the network, each distinct set of virtual ports is assigned a respective distinct broadcast address. The method comprises assigning a respective egress address to each packet entering the network via an ingress virtual port. The respective egress address corresponds to a respective destination address of the entering packet when

a correspondence between the destination address and an egress address is known. When no correspondence between the destination address and an egress address is known, the respective egress address is a broadcast egress address corresponding to the set comprising the ingress virtual port. The method further comprises routing the packet according to the respective egress address. The routing is restricted to virtual ports belonging to the distinct set of virtual ports which includes the ingress virtual port.

The distinct sets of virtual ports and their associated distinct broadcast addresses define isolated virtual private networks within the network. Because the number of different broadcast addresses is much greater than the number of different VLAN identifiers permitted under the IEEE 802.1 standard, the communications network can provide a larger number of isolated virtual private networks than can a standard IEEE 802.1 VLAN network.

Each physical port of the network may map one-to-one onto a corresponding virtual port, or may map onto a corresponding plurality of virtual ports. In the case that a physical port maps onto a plurality of virtual ports, each virtual port of the plurality is associated with a respective distinct combination of a physical address of the physical port and a respective virtual network identifier.

The invention enables network providers and their multiple customers to ensure that data cannot be sent between virtual ports belonging to different distinct sets of virtual ports. Consequently, data sent into a network of virtual ports via one of the virtual ports (the ingress virtual port for that data) can exit the network only at a virtual port (the egress virtual port for that data) belonging to the same distinct set as the ingress port. This property allows the network providers and their multiple customers to ensure that communications between customers can occur only in controlled ways.

This property of the invention may be exploited by arranging that each distinct set of virtual ports is in the control of a single organization. In the case that one and only one virtual port maps to one physical port, 5 the physical port is further arranged to be in the control of the organization that controls the virtual port.

If each virtual port of a particular distinct set of virtual ports is thus mapped to a distinct a physical port, and if no other virtual ports are mapped to those 10 physical ports, than an organization that controls all the virtual ports of the particular set of virtual ports can be assured that only data that originates at one or more of its physical ports can be received at any of its physical ports.

In the case that multiple organizations have 15 elected to trust a service provider to respect their security requirements, multiple virtual ports, each belonging to a different distinct set of virtual ports belonging to a different organization, can be mapped to a 20 physical port belonging to the trusted service provider. The trusted service provider is thereby enabled to communicate with multiple customers through a single physical port, a much more economical arrangement than requiring the service provider to have a separate physical 25 port for each customer.

When the destination address of the packet is a unicast address and a correspondence between the destination address and a unicast egress address is known, the step of assigning an egress address may comprise 30 assigning the unicast egress address. The unicast egress address corresponds to an egress virtual port belonging to the distinct set of virtual ports which includes the ingress virtual port. The destination address is accessible from that egress virtual port. The step of 35 routing the packet may comprise routing the packet to that egress virtual port.

When the destination address of the packet is a unicast address and no correspondence between the destination address and an egress address is known, the step of assigning an egress address may comprise assigning  
5 a broadcast egress address corresponding to the distinct set of virtual ports which includes the ingress virtual port. The step of routing the packet may comprise routing the packet to each virtual port, other than the ingress virtual port, of the distinct set of virtual ports which  
10 includes the ingress virtual port.

When the destination address of the packet is a multicast address, the step of assigning an egress address may comprise assigning a broadcast egress address  
15 corresponding to the distinct set of virtual ports which includes the ingress virtual port. The step of routing the packet may comprise routing the packet to each virtual port of the distinct set of virtual ports which includes the ingress virtual port, other than the ingress virtual port.

20 Alternatively, when the destination address of the packet is a multicast address and a correspondence between the destination address and a multicast egress address is known, the step of assigning an egress address may comprise assigning the multicast egress address. The  
25 multicast egress address corresponds to a plurality of virtual ports belonging to the distinct set of virtual ports which includes the ingress virtual port. The step of routing the packet may comprise routing the packet to each virtual port of the plurality of virtual ports  
30 belonging to the distinct set of virtual ports which includes the ingress virtual port.

The method may further comprise assigning a respective ingress address to each packet entering the network, the respective ingress address corresponding to a  
35 virtual port at which the packet enters the network. The assigned ingress addresses may be used to populate address association tables, and the address association tables may

be used to determine correspondences between destination addresses and egress addresses.

The egress address assigned to a packet may be encapsulated in the packet at the ingress virtual port via which the packet enters the network, and may be removed from the encapsulated packet at an egress virtual port where the packet leaves the network.

A respective ingress address may also be assigned to each packet entering the network, the respective ingress address corresponding to the ingress virtual port via which the packet enters the network. The assigned ingress address may also be encapsulated in the packet as it enters the network. An address association table associated with each virtual port of the network may be maintained, each address association table mapping each of a plurality of egress addresses to at least one corresponding destination address. The address association tables may be used to determine correspondences between destination addresses and egress addresses. On receipt of a packet entering the network via an ingress virtual port, an entry is added to the address association table associated with the ingress virtual port when the address association table does not contain a source address of the packet in any destination address field of the address association table. The entry comprises the source address in a destination address field and the ingress address in a corresponding egress address field. On receipt of an encapsulated packet at a virtual port of the network, an entry is added to the address association table associated with said virtual port when the address association table does not contain a source address of the encapsulated packet in any destination address field of the address association table. The entry comprises the source address in a destination address field and the ingress address of the encapsulated packet in a corresponding egress address field.

The above procedures populate address association tables of the network in a manner that preserves isolation between the communications of distinct customers even though the facilities of the communications network are shared. Consequently, each customer has its own virtual private network provided by the shared facilities.

The routing of packets having broadcast egress addresses may be restricted to only those trunks of the network required to reach virtual ports in the distinct set of virtual ports corresponding to the broadcast egress address. This avoids unwarranted consumption of network resources.

Similarly, the routing of packets having multicast egress addresses may be restricted to only those trunks of network required to reach virtual ports in plurality of virtual ports within a distinct set of virtual ports, the plurality of virtual ports corresponding to the multicast egress address.

Another aspect of the invention provides a communications network comprising a plurality of distinct sets of virtual ports, at least one address assigner and at least one router. No virtual port belongs to more than one of the distinct sets, and each distinct set is assigned a respective distinct broadcast address. Each address assigner is operable to assign a respective egress address to each packet entering the network via an ingress virtual port. The respective egress address corresponds to a respective destination address of the entering packet when a correspondence between the destination address and an egress address is known. The respective egress address is a broadcast egress address corresponding to the set comprising the ingress virtual port when no correspondence between the destination address and an egress address is known. Each router is operable to route the packet according to the respective egress address. The routing is restricted to virtual ports belonging to the distinct



set of virtual ports which includes the ingress virtual port.

As noted above, each physical port of the network may map one-to-one onto a corresponding virtual port, or  
5 may map onto a corresponding plurality of virtual ports. In the case that a physical port maps onto a plurality of virtual ports, each virtual port of the plurality is associated with a respective distinct combination of a physical address of the physical port and a respective  
10 virtual network identifier.

The network may further comprise a plurality of trunks interconnecting routers of the network. Each router is operable to route the packet via trunks of the network. When the packet is assigned a broadcast egress  
15 address corresponding to a distinct set of virtual ports, each router is operable to route the packet via a restricted set of trunks containing only those trunks required to reach virtual ports in the distinct set of virtual ports corresponding to said broadcast egress  
20 address. When the packet is assigned a multicast egress address corresponding to a plurality of virtual ports in a distinct set of virtual ports, each router is operable to route the packet via a restricted set of trunks containing only those trunks required to reach virtual ports in the  
25 plurality of virtual ports corresponding to said multicast egress address.

Yet another aspect of the invention provides a routing device for a communications network. The routing device comprises a plurality of distinct subsets of  
30 virtual ports, at least one address assigner and at least one router. No virtual port belongs to more than one of the distinct subsets. Each distinct subset may be a subset of a respective distinct set of virtual ports of the network. Each distinct set of virtual ports is  
35 assigned a respective distinct broadcast address. Each address assigner is operable to assign a respective egress address to each packet entering the network via an ingress

virtual port of the routing device. The respective egress address corresponds to a respective destination address of the entering packet when a correspondence between the destination address and an egress address is known. The  
5 respective egress address is a broadcast egress address corresponding to the set comprising the ingress virtual port when no correspondence between the destination address and an egress address is known. Each router is operable to route the packet according to the respective  
10 egress address, the routing being restricted to virtual ports belonging to the distinct set of virtual ports which includes the ingress virtual port.

Each router may provide IEEE 802.1 switching functionality adapted to packets encapsulated with ingress  
15 and egress addresses.

A respective address assigner may be provided for each distinct subset of virtual ports. Each address assigner may be connected between its respective distinct subset of virtual ports and a router of the routing  
20 device. The routing device may further comprise a switching element connected between at least one address assigner and its respective distinct subset of virtual ports. The switching element may be operable to multiplex the virtual ports of the respective distinct subset of  
25 virtual ports onto the address assigner. The switching elements may provide IEEE 802.1 switching functionality.

Use of IEEE 802.1 switching functionality enables a NSP to provide transparent Ethernet LAN service across the NSP's network. Transparent Ethernet LAN  
30 service is attractive to many customers, as they are already familiar with the operation of Ethernet networks. Moreover, the use of many Ethernet conventions in the NSP network enable considerable re-use of proven and cost-effective Ethernet hardware and software in constructing  
35 the NSP network, and familiarity with the operation of Ethernet networks will facilitate operation of the shared network by the NSP.

The routing device may further comprise a VLAN demultiplexer connected between the router and a plurality of the address assigners. The VLAN demultiplexer is operable to route an encapsulated packet from the router to an address assigner selected according to the ingress address and the egress address of the encapsulated packet. The routing is such that all encapsulated packets having a common egress address and an ingress address corresponding to a virtual port in a particular set of the distinct sets of virtual ports are routed to an address assigner associated with that egress address and that particular distinct set of virtual ports.

Use of the VLAN demultiplexer permits some sharing of egress addresses among distinct virtual private networks without compromising the isolation between distinct virtual private networks. This capability is useful for connections between the network and external routers (e.g. Internet routers) where a respective dedicated link for each virtual private network is not economically feasible. Where the VLAN demultiplexer is used, a plurality of virtual ports may be connected to a common physical port of the routing device. Each such virtual port is associated with a unique combination of the physical address of the common physical port and a virtual network identifier.

Some translation of virtual private network identifiers may also be provided at interfaces to other networks supporting the virtual private networks.

### Brief Description of Drawings

Embodiments of the invention are described below by way of example only. Reference is made to accompanying drawings, in which:

Figure 1 is a block schematic diagram of a NSP network according to an embodiment of the invention;

Figure 2 is a block schematic diagram of an access switch of the network of Figure 1;

Figure 3 is flow chart illustrating operation of an encapsulation/decapsulation device of the access switch of Figure 1 on receipt of a data frame at a customer port of the access switch;

5           Figure 4 is a flow chart illustrating operation of a multiplex switch of the access switch of Figure 2 on receipt of an encapsulated data frame from the encapsulation/decapsulation device;

10           Figure 5 is a flow chart illustrating operation of the multiplex switch of the access switch of Figure 2 on receipt of an encapsulated data frame from another switch on a trunk;

15           Figure 6 is a flow chart illustrating operation of the encapsulation/decapsulation device on receipt of an encapsulated data frame from the multiplex switch;

            Figure 7 is a block schematic diagram showing a first embodiment 22 of an access switch adapted to support connection of the NSP network to ISP routers;

20           Figure 8 is a flow chart illustrating aspects of the operation of a VLAN demultiplexer of the access switch of Figure 7;

            Figure 9 is a block schematic diagram showing a second embodiment 42 of an access switch adapted to support connection of the NSP network to ISP routers; and

25           Figure 10 is a block schematic diagram showing a third embodiment of an access switch 62 adapted to support connection of the NSP network to ISP routers.

#### Detailed Description of Embodiments

30           Figure 1 is a block schematic diagram of a NSP network 10 according to an embodiment of the invention. The NSP network 10 comprises a plurality of routing devices in the form of access switches 12 interconnected via transmission facilities 14. In some implementations,  
35           one or more core switches 16 may be connected between some of the access switches 12. The access switches 12 are

each connected to one or more customer LANs 20 via respective access links 22.

Figure 2 is a block schematic diagram of an access switch 12 of the network of Figure 1 according to a first embodiment of the invention. The access switch 12 comprises a plurality of address assigners in the form of Encapsulation/Decapsulation Devices (EDDs) 120, each of which is connected to one or more customer ports 123 of the access switch 12 via a respective virtual customer access switch 124. All customer ports 123 associated with a particular EDD 120 and its customer access switch 124 are connected to the same customer LAN 20 via one or more access links 22 - i.e. no customer access switch 124 or EDD 120 has customer ports 123 connected to the customer LANs 20 of more than one customer. The physical customer ports 123 map one-to-one onto respective virtual ports 122. Each customer access switch 124 uses IEEE 802.1 protocols to communicate with the customer LAN 20 to which it is connected via the customer port(s) 123.

The EDDs 120 are also connected to trunks 126 of the access switch 12 via a router in the form of a virtual multiplex switch 127 which operates according to IEEE 802.1D/Q protocols adapted to handle a longer than standard data frame as will be explained below.

Each EDD 120 maintains a respective Destination Address Association Table (DAAT) which maps Medium Access Control (MAC) addresses of elements of the customer LANs 20 in Destination Address (DA) fields onto corresponding customer port addresses in Decapsulation Egress Address (DEA) fields. Each DA is mapped onto a single DEA, but each DEA may be mapped onto a plurality of DAs. Each customer has a unique set of DEAs corresponding to the virtual ports 122 and the associated customer ports 123 connected to that customer's private networks. If distinct customers use the same DA, that DA will be mapped onto a different DEA in the distinct DAATs used for those customers.

A typical customer will have customer LANs 20 using IEEE 802.1 protocols at more than one site and will want to exchange data packets in the form of IEEE 802.3 data frames between elements of the LANs 20 at different sites. As will be explained below, such customers may subscribe to a Carrier Virtual LAN (CVLAN) service provided by the NSP using the NSP network 10. The CVLAN service provides transparent LAN connectivity between customer LANs at different sites with full isolation between the virtual private LANs (or CVLANs) of many distinct customers.

An IEEE 802.3 data frame has a header comprising a Destination Address (DA) identifying a LAN element for which the data frame is intended and a Source Address (SA) identifying the LAN element from which the data frame is sent. When an IEEE 802.3 data frame addressed to a DA on a customer's LAN 20 at a one site is sent on that customer's LAN 20 at another site, the customer's LAN 20 at the other site will route the frame to an access switch 12 connected to the customer's LAN 20 at the other site.

The access switch 12 receives the frame via a customer port 123 connected to the customer LAN 20 at the other site and routes the frame via the associated virtual port 122 and the customer access switch 124 to the EDD 120 for that customer at that access switch 12.

Figure 3 is flow chart illustrating operation of the EDD 120 on receipt of the data frame via the customer port 122. The EDD 120 searches its DAAT for the DA of the received frame. If the DA is in the DAAT, the EDD 120 reads the DEA corresponding to the DA from the DAAT 129. If the DEA corresponds to the customer port 123 on which the frame was received, the frame is intended for an element of the customer's LAN 20 on which the frame was sent. In this case, the EDD 120 discards the frame since no transmission of the frame across the NSP network 10 is required.

However, if the DEA is not equal to the address of the customer port 123 on which the frame was received, the frame is intended for the customer's LAN 20 at another site. In this case, the frame is encapsulated by adding  
5 an additional header that includes the DEA and an Encapsulation Ingress Address (EIA) set equal to the address of the customer port 123 on which the frame was received. As will be explained below, the DEA is used to route the encapsulated frame through the NSP network 10 to  
10 a virtual port 122 and its associated customer port 123. The customer port 123 has an address corresponding to the DEA, and is connected to the customer LAN 20 on which the DA will be found.

If the DA is not found in the DAAT, the EDD 120  
15 is unable to map the DA onto a corresponding DEA to route the frame across the NSP network 10. In this case, the EDD 120 encapsulates the frame with the DEA set to a CVLAN Broadcast Address (CBA) which enables the frame to be routed to all access switches 12 serving the CVLAN.  
20 Because the EDD serves only a single customer, the CBA can be made specific to that customer so that the frame is routed only to virtual ports 122 and associated customer ports 123 connected to sites of that customer.

If the DA of the received frame is a multicast  
25 address, the EDD 120 sets the DEA equal to a multicast egress address. This multicast egress address may correspond to the CBA of the CVLAN if multiple multicast groups within the CVLAN are not supported, or may correspond to a multicast address that is particular to  
30 the multicast group within the CVLAN if multiple multicast groups with the CVLAN are supported. Such egress address assignments may be arranged through suitable entries in the DAAT or by other means.

Unnecessary broadcasting of frames in the NSP  
35 network 10 wastes network resources. Consequently, the EDD 120 assesses whether the received frame contains information that can be used to augment the DAAT 129. In

particular, when a frame having a particular network address in the SA field is received on a particular customer port 122, it can be inferred that this particular network address can be accessed via this particular customer port 122. Consequently, there should be an entry in the DAAT mapping the network address in the SA field onto the network address of the customer port 122.

The EDD determines whether that entry is missing from the DAAT by searching for the SA of the received frame in the DA fields of the DAAT. If the SA is found, the entry already exists. However, if the SA is not found, the EDD adds an entry to the DAAT, the entry having the SA of the received frame in the DA field and the address of the customer port 122 in the DEA field.

In addition to encapsulating the frame with the EIA and the DEA, the EDD 120 may encapsulate the frame with an Encapsulating VLAN tag (EVTAG) field similar to the VLAN tag of a standard IEEE 802.3 frame. The EVTAG field may contain a 12 bit VLAN identifier and a 3 bit Quality of Service (QoS) indicator.

The frame may also be encapsulated with a Header Checksum, a 32 bit value that will produce an all 1's value in a Cyclic Redundancy Check (CRC) register when a standard IEEE 802.3 checksum CRC procedure is applied to the encapsulation header including the Header Checksum. The all 1's value is the normal starting value for the CRC register in the IEEE 802.3 checksum procedure. The presence of this value in the CRC register at the end of the Header Checksum means that the IEEE 802.3 Checksum field, that was calculated and appended to the unencapsulated frame when the unencapsulated frame was created, can be used unchanged to protect the whole encapsulated frame during transmission through the NSP network 10. Consequently, IEEE 802.1D bridging can be used to forward encapsulated frames, provided only that the multiplex switches 127 are adapted to handle frames longer than standard IEEE 802.3 frames while preserving



and using the Checksum values calculated at creation of the unencapsulated frames.

Figure 4 is a flow chart illustrating operation of the multiplex switch 127 on receipt of an encapsulated frame from the EDD 120. The multiplex switch 127 is similar to a IEEE 802.1D/Q switch adapted to handle the increased length of the encapsulated frame and to operate on the added header.

On receipt of an encapsulated frame, the multiplex switch 127 reads the DEA from the header of the encapsulated frame and determines whether the DEA is a CBA. If the DEA is not a CBA, the multiplex switch 127 finds trunk 126 corresponding to the DEA in a routing table and forwards the encapsulated frame to that trunk 126. If the DEA is a CBA or a multicast egress address, the multiplex switch determines which trunks are registered for that CBA and forwards the encapsulated frame to all trunks 126 registered for that CBA. (The process of trunk registration is described in greater detail below.)

Any core switches 16 in the NSP network 10 operate essentially as described above for the multiplex switch 127 on receipt of an encapsulated frame at a trunk of the core switch 16.

Figure 5 is a flow chart illustrating operation of the multiplex switch 127 on receipt of an encapsulated data frame from another switch of the NSP network 10 on a trunk 126 of the multiplex switch 127. The multiplex switch 127 reads the DEA from the header of the encapsulated frame. If the DEA is not a CBA, the multiplex switch 127 finds the EDD 120 corresponding to the DEA in a routing table and forwards the encapsulated frame to that EDD 120. If the DEA is a CBA, the multiplex switch 127 finds all EDDs 120 corresponding to the CBA and floods the encapsulated frame to all EDDs 120 corresponding to the CBA.

Figure 6 is a flow chart illustrating operation of the EDD 120 on receipt of an encapsulated data frame from the multiplex switch 127. The EDD 120 reads the DEA from the encapsulated frame and compares the DEA to the addresses of the customer ports 122 connected to the EDD 120 via the customer access switch 124. If the DEA matches the address of a customer port 123 connected to the EDD 120, the EDD 120 decapsulates the frame by removing the header containing the DEA and the EIA, and routes the decapsulated frame to the customer port 123 via the customer access switch 124 and the virtual port 122.

If the DEA does not match the address of any customer port 123 connected to the EDD 120, the EDD 120 determines whether the DEA is a CBA for the EDD 120. If the DEA is a CBA for the EDD 120, the EDD 120 decapsulates the frame by removing the header containing the DEA and the EIA, and routes the decapsulated frame to all customer ports 123 corresponding to the CBA.

If the DEA does not match the address of any customer port 123 connected to the EDD 120 and is not a CBA for the EDD 120, the frame is not forwarded to any customer port 123.

The EDD 120 also assesses whether the received encapsulated frame contains information that can be used to augment the DAAT. In particular, the EDD 120 searches for the SA of the received encapsulated frame in the DA fields of the DAAT. If the SA is found, the entry already exists. However, if the SA is not found, the EDD adds an entry to the DAAT, the entry having the SA of the received frame in the DA field and the EIA of the encapsulated frame in the DEA field.

It follows from the operations of the elements of the NSP network 10 as described above, that a typical IEEE 802.3 frame is routed across the NSP network 10 from a first site of a customer LAN 20 to a second site of the customer LAN 20 as follows:

1. The IEEE 802.1 frame is routed by the customer LAN 20 at the first site to a first access switch 12 serving the first site based on the DA of the frame.
2. The IEEE 802.3 frame is encapsulated at the first  
5 access switch 12 by adding a header comprising a DEA specifying a port on a second access switch 12 serving the second site of the customer LAN 20.
3. The encapsulated frame is routed across the NSP network  
10 from the first access switch 12 to the second access switch 12 based on the DEA of the encapsulated frame.
4. The encapsulated frame is decapsulated by the second access switch 12 and forwarded to the second site of the customer LAN where it is routed based on the DA of the decapsulated frame.

15           When the access switch 12 receiving the frame from the first site of the customer LAN 20 is unable to determine the DEA from the DA of the received frame, the frame is flooded across the network to all sites of the customer LAN 20 as follows:

- 20 1. The IEEE 802.3 frame is routed by the customer LAN 20 at the first site to a first access switch 12 serving the first site based on the DA of the frame.
2. The IEEE 802.3 frame is encapsulated at the first  
access switch 12 by adding a header comprising a CBA in  
25 the DEA field.
3. The encapsulated frame is flooded across the NSP network 10 from the first access switch 12 to all access switches 12 serving sites of the customer LAN 20 based on the CBA of the encapsulated frame.
- 30 4. The encapsulated frame is decapsulated by the destination access switches 12 and forwarded to the other sites of the customer LAN where it is routed based on the DA of the decapsulated frame.

35           Similarly, IEEE 802.3 frames having a multicast address in the DA field are encapsulated with the CBA in the DEA field and are flooded across the NSP network 10

from the first access switch to all access switches 12  
serving sites of the customer LAN 20.

The DEAs used for a particular customer are  
unique to that customer because of the technique used to  
5 fill the DAAT at each EDD 120. Each EDD 120 is assigned  
to a single customer and serves only virtual ports 122 and  
associated customer ports 123 which are assigned that  
customer. When an EDD 120 adds an entry to its DAAT based  
on receipt of an unencapsulated frame from a connected  
10 customer port 122, the DEA of that entry must be the DEA  
of the customer port 122 which is uniquely assigned to  
that customer. When an EDD 120 receives an encapsulated  
frame from the multiplex switch 127, it verifies that the  
frame has a DEA corresponding to a connected customer port  
15 123 or a CBA corresponding to its assigned customer to  
ensure that the frame comes from within the CVLAN of its  
customer before adding any entry to its DAAT. Such an  
entry must include the EIA of the frame in the DEA field,  
and that EIA corresponds to a customer port 122 that is  
20 assigned to the same customer - otherwise the received  
frame would not have a DEA or CBA corresponding to that  
customer.

Because the virtual ports 122 and associated  
physical customer ports 123 connected to each customer LAN  
25 20 and the corresponding EDDs 120, DAATs, DEAs and CBAs  
are unique to that particular customer, frames cannot be  
transmitted from one customer to any other customer even  
though the frames are transmitted over a shared NSP  
network 10. Consequently, each customer has a CVLAN that  
30 is isolated from the CVLANs of other customers. The NSP  
network 10 can provide a very large number of isolated  
CVLANs to serve a very large number of customers because  
the isolation between CVLANs is determined by unique sets  
of virtual ports and associated broadcast addresses rather  
35 than by a more limited number of CVLAN identifiers.

However, only the virtual ports 122 and  
associated customer ports 123, the customer access

switches 124, the EDDs 120 and the DAATs are dedicated to specific customers. The multiplex switches 127, core switches 16 and transmission facilities 14 are shared among many customers for economies of scale. Moreover, 5 key elements of the customer access switches 124, multiplex switches 127 and core switches 16 can be provided using proven IEEE 802.1D/Q hardware and software with relatively minor modifications for further cost advantages. The extensive use of modified IEEE 802.1D/Q 10 techniques in this embodiment of the NSP network 10, also ensures that extensive industry experience in operating IEEE 802.1 networks can be applied readily to the operation of this network.

The above description refers to registration of 15 CBAs at trunks 126 of the access switches 12. IEEE 802.1D defines procedures for registering multicast groups at trunks such that frames carrying a particular multicast address in the DA field are forwarded only by trunks which have that multicast address registered for that trunk. 20 The multicast group registrations are propagated by the IEEE 802.1D GARP Multicast Registration Protocol (GMRP) to all trunks in the network needed to create a minimal subset of interconnections that interconnects all registrants to the group.

25 These multicast group registration techniques can be adapted to the registration of trunks for CBAs in the NSP network 10. Each EDD 120 registers a corresponding CBA at its multiplex switch port so that encapsulated frames having a particular CBA in the DEA field will be 30 transmitted over only those trunks needed to transmit the frame to the other EDDs 120 of the particular CVLAN corresponding to the CBA. This avoids wasteful transmission of frames to EDDs 120 that are not participating in the CVLAN.

35 According to the description given above, all frames having a multicast DA may be assigned a selected CBA for a DEA, the CBA being selected according to the

ingress port at which the frame was received. While this procedure restricts frames to the CVLANs for which they are intended, it does not enable customers to restrict multicast frames to distinct multicast groups within their  
5 CVLANs.

Distinct multicast groups within CVLANs can be supported by defining a distinct multicast DEA for each such multicast group. The multicast DEAs must be unique to the CVLAN to which the multicast group belongs, and the  
10 EDDs 120 must translate multicast DAs of unencapsulated frames entering the NSP network 10 into the appropriate multicast DEAs using the DAATs or some other means. The multicast DEAs should be locally administered by the NSP.

The NSP can ensure that each multicast DEA is  
15 unique to a particular CVLAN within the NSP network 10 is by requiring a multicast DEA format that combines a CVLAN identifier with a multicast group identifier. For example, each multicast DEA could comprise:

1. a multicast bit (indicating whether the address is a  
20 unicast address or a multicast address),
2. a local administration bit (indicating whether the address is locally administered),
3. a CVLAN identifier (identifying the CVLAN to which the packet is to be restricted),
- 25 4. an IP multicast bit (indicating whether the multicast is an IP multicast), and
5. a multicast group identifier (identifying the multicast group within the CVLAN to which the packet is to be restricted).

30 The local administration bit can be used to detect frames bearing multicast addresses that are not locally administered so that such frames can be discarded to ensure that isolation between distinct CVLANs is preserved.

35 The multicast group identifier can be the multicast DA or an identifier derived from the multicast DA. Because the multicast DEAs include a CVLAN

identifier, the same multicast DAS can be used in distinct CVLANs without loss of isolation between distinct CVLANs.

According to this addressing scheme, the CBA for a particular CVLAN could comprise:

- 5 1. a 1 for the multicast bit,
2. a 1 for the local administration bit,
3. the CVLAN identifier for the particular CVLAN'
4. a 0 for the IP multicast bit, and
5. a field of 0's for the multicast group identifier.

10       The IEEE 802.1D GARP Multicast Registration Protocol (GMRP) referenced above can be modified for NSP networks 10 supporting multicast groups within CVLANs to create a minimal subset of interconnections that interconnects all registrants to the multicast group. In  
15 particular, the GMRP is modified to ensure that GMRP messages related to multicast DEAs other than CBAs are transmitted and create trunk registrations only on trunks registered for the CBA of the CVLAN to which the multicast DEA belongs. Consequently, GMRP message activity for  
20 multicast DEAs other than CBAs are confined to the physical topology in which messages addressed by the CBA can propagate. GMRP messages required for the registration of CBAs are not so confined, but such messages are infrequent because new registrations for CBAs  
25 occur only when a new customer site is configured.

A frame bearing a multicast DEA other than a CBA may be transmitted on a trunk only if the trunk has received a GMRP group registration generated by a GMRP application from another switch. This is the fundamental  
30 multicast tree pruning rule of IEEE 802.1D "extended filtering". This technique achieves bandwidth savings by ensuring that multicast frames are transmitted on trunks only if a station that can be reached on that trunk has indicated an interest in receiving multicasts from that  
35 multicast group.

EDDs 120 of the NSP network 10, must translate IGMP join requests entering the NSP network 10 into GMRP

join requests for forwarding into the NSP network 10 according to the modified GMRP procedures described above.

In the NSP network 10 described above, each CVLAN is defined by a distinct set of virtual ports 122 have a one-to-one mapping to a respective distinct set of customer ports 123 having physical addresses defining a distinct set of respective egress addresses. According to this scheme for isolating distinct CVLANs in the NSP network 10, each CVLAN would require a separate physical port and transmission link for connection to each ISP router to which connection of the CVLAN is required. However, it is not economically feasible to provide a separate dedicated link for connection of each CVLAN to each ISP router. Consequently, alternative arrangements are required for connection of the NSP network 10 to ISP routers over transmission links shared among CVLANs. The alternative arrangements must preserve the isolation between the CVLANs.

Figure 7 is a block schematic diagram showing a first embodiment 22 of an access switch adapted to support connection of the NSP network 10 to ISP routers 300, 302. The ISP routers 300, 302 are IEEE 802.1 routers that use VLAN tags to separate CVLANs.

The access switch 22 comprises a plurality of address assigners in the form of EDDs 120 and a router in the form of a virtual multiplex switch 127 as did the access switch 12. The access switch 22 further comprises a plurality of VLAN demultiplexers 222 connected between the multiplex switch 127 and groups of the EDDs 120, each VLAN demultiplexer 222 being associated with a respective egress address or a respective distinct set of egress addresses. Each EDD 120 is connected to a respective virtual port 122. A respective VLAN translator 224 is connected to each virtual port 122, and each group of VLAN translators 224 is connected to a respective router demultiplexer 226. The router demultiplexers 226 are connected to external ISP routers 300, 302.



On receipt of an encapsulated packet having an egress address corresponding to one of the external routers 300, 302 via a trunk 126, the virtual multiplex switch 127 routes the encapsulated packet to a VLAN demultiplexer 222 selected according to the egress address. The selected VLAN demultiplexer 222 routes the encapsulated packet to an EDD 120 selected according to the ingress address of the encapsulated packet. This selection scheme ensures that all encapsulated packets having a common egress address and an ingress address corresponding to a virtual port 122 in a particular set of the distinct sets of virtual ports 122 are routed to an EDD 120 associated with that egress address and that particular distinct set of virtual ports 122.

Because the egress address of a packet directed to an ISP router 300, 302 identifies the ISP router 300, 302, it does not uniquely identify the CVLAN to which the packet is to be restricted. Consequently, the VLAN demultiplexer 222, uses the ingress address of the packet to determine which EDD 120 should process the packet since the ingress address does uniquely identify the CVLAN to which the packet is restricted. However, when the egress address is a broadcast or multicast egress address employing the format described above for broadcast and multicast egress addresses, the VLAN demultiplexer 222 may determine which EDD 120 to route the packet to, either from the egress address or from the ingress address.

Each VLAN demultiplexer 222 may maintain a table for associating ingress addresses with EDDs 120 and may employ that table to determine the routing of packets to EDDs 120. The VLAN demultiplexers 222 may use the ingress addresses and egress addresses of broadcast and multicast packets to populate the table. In particular, when a VLAN demultiplexer 222 receives a broadcast or multicast packet having an ingress address that does not appear in any ingress address field of the table, it may create a new entry having the ingress address in an ingress address

field of the table and an EDD identifier determined from the broadcast or multicast egress address of the packet.

Figure 8 is a flow chart illustrating operation of the VLAN demultiplexers 222 on receipt of a packet from the multiplex switch 127 in more detail.

The selected EDD 120 decapsulates the packet and forwards it via the respective virtual port 122 to the respective VLAN translator 224. The VLAN translator 224 applies a respective VLAN identifier to the packet. The VLAN identifier corresponds to the distinct set of ports containing the ingress port, i.e. it is particular to the CVLAN which corresponds to that distinct set of ports. The VLAN translator 224 forwards the resulting packet to the router demultiplexer 226.

The VLAN translators 224 may receive broadcast packets for VLANs to which are not supported by the ISP routers 300, 302. The VLAN translators 224 discard such packets.

The router demultiplexer 226 routes the packet to an IEEE 802.1 external router 300. The external router 300 preserves isolation of CVLANs using VLAN identifiers according to the IEEE 802.1 standard.

On receipt of a packet from one of the external routers 300, the router demultiplexer 226 routes the packet to VLAN translator 224 selected according to a VLAN identifier of the received packet. The VLAN translator 224 forwards the packet to its respective EDD 120. The EDD 120 encapsulates the packet with an ingress address corresponding to its respective virtual port 122 and an egress address corresponding to its destination address, and forwards the encapsulated packet to the VLAN demultiplexer 222. The VLAN demultiplexer 222 forwards the encapsulated packet to the virtual multiplex switch 127 for routing according to the egress address.

Note that the arrangement described above enables a particular CVLAN within the network 10 to be mapped onto one VLAN identifier in a first IEEE 802.1 VLAN identifier

space supported by a first external router 300 or plurality of routers 300. The same CVLAN within the network 10 may be mapped onto another VLAN identifier in a second IEEE 802.1 VLAN identifier space supported by a second external router 302 or plurality of routers 302, so assignment of VLAN identifiers in distinct external IEEE 802.1 VLAN networks need not be coordinated. Moreover, the arrangement described above enables the same VLAN identifier in different IEEE 802.1 VLAN identifier spaces to be mapped onto different CVLANs in the network 10. This is advantageous because, as noted above, each IEEE 802.1 VLAN identifier space is limited to 4095 distinct VLANs, whereas the network 10 can support many times that number of CVLANs.

In the embodiment of Figure 7, the virtual ports 122 have the same properties as the virtual ports 122 of the embodiment of Figure 2. In particular, each CVLAN has a distinct set of virtual ports 122, no virtual port 122 belonging to more than one of the distinct sets.

In the arrangement of Figure 7, each customer can choose his router-access VLAN identifiers arbitrarily. There is no requirement that VLAN identifier choice be coordinated between multiple customers. Each ISP router 300, 302 participates in only one VLAN identifier space. The access switch 22 translates VLAN identifiers between this one VLAN identifier space and the many VLAN identifier spaces of the NSP network 10. The NSP network 10 has one VLAN identifier space for each distinct CVLAN. Each ISP router 300, 302 may either share a VLAN identifier space with one or more other routers belonging to the same ISP or have its own dedicated VLAN identifier space.

The NSP must establish an association between each customer VLAN requiring ISP router access and a unique VLAN in each ISP router VLAN identifier space. This association requires a three-way agreement between the customer, the NSP and the ISP, as follows:

1. The ISP needs to know, for each customer, which subnets are to be supported. The NSP decides which of his VLAN identifiers he will assign to each subnet.

2. Each customer needs to know the subnet mask and router IP address for each subnet and which of his VLAN identifiers he will assign to each subnet.

3. The NSP needs to know the pairing of VLANs created by the decisions taken by the ISP and the customer to support the subnet. The VLAN pairing created for each subnet must be configured in the VLAN translating access switch 22 so that VLAN identifiers may be modified in packets passing between router access VLAN identifier spaces and customer VLAN identifier spaces.

Figure 9 is a block schematic diagram showing a second embodiment 42 of an access switch adapted to support connection of the network 10 to ISP routers 500, 502. The ISP routers 500, 502 are MPLS routers providing multiple virtual router capability.

The access switch 42 comprises a plurality of address assigners in the form of EDDs 120 and a router in the form of a virtual multiplex switch 127 as did the access switches 12, 22. The access switch 42 further comprises a plurality of VLAN demultiplexers 222 connected between the multiplex switch 127 and groups of the EDDs 120, each VLAN demultiplexer 222 being associated with a respective egress address as in the access switch 22. Each EDD 120 is connected to a respective virtual port 122. A respective Multi-Protocol Label Switching (MPLS) converter 424 is connected to each virtual port 122, and the MPLS converters 424 are connected to a MPLS switch 426.

On receipt of an encapsulated packet on a trunk 126, the virtual multiplex switch 127 routes the encapsulated packet to a VLAN demultiplexer 222 selected according to the egress address. The selected VLAN demultiplexer 222 routes the encapsulated packet to an EDD 120 selected according to the ingress address of the

encapsulated packet. This selection scheme ensures that all encapsulated packets having a common egress address and an ingress address corresponding to a virtual port 122 in a particular set of the distinct sets of virtual ports 122 are routed to an EDD 120 associated with that egress address and that particular distinct set of virtual ports 122.

The selected EDD 120 decapsulates the packet and forwards it via the respective virtual port 122 to the respective MPLS converter 424. The MPLS converter 424 applies a respective MPLS label to the packet. The MPLS label corresponds to the distinct set of virtual ports 122 containing the ingress virtual port 122, i.e. it is particular to the CVLAN which corresponds to that distinct set of virtual ports. The MPLS converter 424 forwards the resulting packet to the MPLS switch 426. The MPLS switch 426 routes the packet to an external router 500. The external router 500 preserves isolation of CVLANs using the MPLS labels that are unique to CVLAN.

On receipt of a packet from one of the external routers 500, the MPLS switch 426 routes the packet to a MPLS converter 424 selected according to a MPLS label of the received packet. The MPLS converter 424 forwards the packet to its respective EDD 120 via its respective virtual port 122. The EDD 120 encapsulates the packet with an ingress address corresponding to its respective virtual port 122 and an egress address corresponding to its destination address, and forwards the encapsulated packet to the VLAN demultiplexer 222. The VLAN demultiplexer 222 forwards the encapsulated packet to the virtual multiplex switch 127 for routing according to the egress address.

Note that the arrangement described above enables a particular CVLAN within the network 10 to be mapped onto one MPLS label in a first MPLS label space supported by a first external router 500 or plurality of routers 500. The same CVLAN within the network 10 may be mapped onto

another MPLS label in a second MPLS label space supported by a second external router 502 or plurality of routers 502.

Figure 10 is a block schematic diagram showing a  
5 third embodiment of an access switch 62 adapted to support connection of the network 10 to ISP routers 700.

The access switch 62 comprises a plurality of address assigners in the form of EDDs 120 and a router in the form of a virtual multiplex switch 127 as did the  
10 access switches 12, 22, 42. The access switch 62 further comprises a plurality of VLAN demultiplexers 222 connected between the multiplex switch 127 and groups of the EDDs 120, each VLAN demultiplexer 222 being associated with a respective egress address as in the access switches 22,  
15 42. Each EDD 120 is connected to a respective virtual port 122. A respective virtual private router 624 is connected to each virtual port 122, and each virtual private router 624 is connected to respective network address translator 626.

On receipt of an encapsulated packet on a trunk  
20 126, the virtual multiplex switch 127 routes the encapsulated packet to a VLAN demultiplexer 222 selected according to the egress address. The selected VLAN demultiplexer 222 routes the encapsulated packet to an EDD  
25 120 selected according to the ingress address of the encapsulated packet. This selection scheme ensures that all encapsulated packets having a common egress address and an ingress address corresponding to a virtual port 122 in a particular set of the distinct sets of virtual ports  
30 122 are routed to an EDD 120 associated with that egress address and that particular distinct set of virtual ports 122.

The selected EDD 120 decapsulates the packet and forwards it via the respective virtual port 122 to the  
35 respective virtual private router 624. The virtual private router 624 discards any packets not having a destination IP address corresponding to the router 700

connected to the respective network address translator 626, and forwards any packets having a destination address corresponding to the router 700 to the respective network address translator 626. The network address translator  
5 626 translates the destination address from a private IP address in the customer's private IP address space to a corresponding public IP address in the public IP address space. The network address translator 626 forwards the packet with the translated IP address to the router 700.

10 On receipt of a packet from one of the external routers 700, a network address translator 626 translates the destination address of the received packet from a public IP address to a corresponding private IP address in the private IP address space of the NSP network 10. The  
15 network address translator 626 forwards the packet with the translated IP address to its respective virtual private router 624. The virtual private router 624 applies a corresponding MAC destination address to the packet in the DA field and forwards the resulting packet  
20 to its respective EDD 120 via its respective virtual port 122. The EDD 120 encapsulates the packet with an ingress address corresponding to its respective virtual port 122 and an egress address corresponding to its destination address, and forwards the encapsulated packet to the VLAN  
25 demultiplexer 222. The VLAN demultiplexer 222 forwards the encapsulated packet to the virtual multiplex switch 127 for routing according to the egress address.

Note that the arrangement described above enables a particular CVLAN within the network 10 to be mapped onto  
30 a restricted set of IP addresses in the IP routers 700.

In the arrangement of Figure 10, one or more of the IP routers could be integrated into the access switch 62 to provide an IP router appropriate for direct connection to the NSP network 10.

35 Some or all of the network address translators 626 of Figure 10 could be eliminated if the IP addresses corresponding to one or more of the virtual private

networks in the NSP network are registered as public IP addresses.

Moreover, the arrangements of two or more of Figures 2, 7, 9 and 10 could be integrated into a single  
5 access switch in which a virtual multiplex switch 127 is shared between the combined arrangements. In this case, and in networks that combine the functionality of one or more of Figures 7, 9 and 10 with the functionality of Figure 2, each distinct set of virtual ports 122 defining  
10 a virtual private network may include some virtual ports 122 which map one-to-one onto corresponding physical ports, such as the customer ports 123 of the Figure 2 embodiment. The physical ports are each associated with a unique respective physical address. Other groups of  
15 virtual ports 122 may be connected to a common physical port for each group. Each such virtual port 122 is associated with a unique combination of the physical address of the common physical port and some other identifier that identifies the virtual private network  
20 with which the virtual port 122 is associated. The other identifier may be one or more of an ingress address, a virtual private network identifier, a VLAN identifier, an MPLS label or any other identifier sufficient to unambiguously determine the virtual private network with  
25 which the virtual port 122 is associated.

While embodiments of the invention are described above in terms of standard IEEE 802.3 frames and IEEE 802.1 protocols, the invention could be practised with other frame formats and protocols. While encapsulation  
30 with IEEE 802.1 addresses is described above, the frames could be encapsulated with other types of addresses, such as IP addresses, for example.

These and other variations do not depart from the principles of the invention as defined by the claims  
35 below.



We claim:

1. A method of routing packets through a communications network having a plurality of distinct sets  
5 of virtual ports, no virtual port belonging to more than one of the distinct sets, a respective distinct broadcast address being assigned to each distinct set of virtual ports, the method comprising:

assigning a respective egress address to each  
10 packet entering the network via an ingress virtual port, the respective egress address corresponding to a respective destination address of the entering packet when a correspondence between the destination address and an egress address is known, and the respective egress address  
15 being a broadcast egress address corresponding to the set comprising the ingress virtual port when no correspondence between the destination address and an egress address is known; and

routing the packet according to the respective  
20 egress address, said routing being restricted to virtual ports belonging to the distinct set of virtual ports which includes the ingress virtual port.

2. A method as defined in claim 1, wherein,  
25 when the destination address of the packet is a unicast address and a correspondence between the destination address and a unicast egress address is known:

the step of assigning an egress address comprises  
assigning the unicast egress address, said unicast egress  
30 address corresponding to an egress virtual port belonging to the distinct set of virtual ports which includes the ingress virtual port, the destination address being accessible from said egress virtual port; and

the step of routing the packet comprises routing  
35 the packet to said egress virtual port.

3. A method as defined in claim 1, wherein, when the destination address of the packet is a unicast address and no correspondence between the destination address and an egress address is known:

5 the step of assigning an egress address comprises assigning a broadcast egress address corresponding to the distinct set of virtual ports which includes the ingress virtual port; and

10 the step of routing the packet comprises routing the packet to each virtual port, other than the ingress virtual port, of the distinct set of virtual ports which includes the ingress virtual port.

4. A method as defined in claim 1, wherein, when the destination address of the packet is a multicast address:

15 the step of assigning an egress address comprises assigning a broadcast egress address corresponding to the distinct set of virtual ports which includes the ingress virtual port; and

20 the step of routing the packet comprises routing the packet to each virtual port of the distinct set of virtual ports which includes the ingress virtual port other than the ingress virtual port.

25

5. A method as defined in claim 1, wherein, when the destination address of the packet is a multicast address and a correspondence between the destination address and a multicast egress address is known:

30 the step of assigning an egress address comprises assigning the multicast egress address, said multicast egress address corresponding to a plurality of virtual ports belonging to the distinct set of virtual ports which includes the ingress virtual port; and

35 the step of routing the packet comprises routing the packet to each virtual port of said plurality of

virtual ports belonging to the distinct set of virtual ports which includes the ingress virtual port.

6. A method as defined in claim 1, further comprising:

assigning a respective ingress address to each packet entering the network, the respective ingress address corresponding to a virtual port via which the packet enters the network;

10 using the assigned ingress addresses to populate address association tables; and

using the address association tables to determine correspondences between destination addresses and egress addresses.

15 7. A method as defined in claim 1, further comprising:

adding to each packet entering the network via an ingress virtual port the respective egress address assigned to that packet to provide a corresponding encapsulated packet;

routing the encapsulated packet in the network according to assigned egress address encapsulated in the packet; and

25 removing from each encapsulated packet received at an egress virtual port of the network the egress address assigned to that packet to provide a decapsulated packet.

30 8. A method as defined in claim 7, further comprising:

assigning a respective ingress address to each packet entering the network, the respective ingress address corresponding to the ingress virtual port via which the packet enters the network;

35

adding the assigned ingress address to each packet entering the network in providing the corresponding encapsulated packet;

maintaining an address association table  
5 associated with each virtual port of the network, each address association table mapping each of a plurality of egress addresses to at least one corresponding destination address; and

using the address association tables to determine  
10 correspondences between destination addresses and egress addresses, wherein:

on receipt of a packet entering the network via an ingress virtual port, said packet including a source  
15 address, an entry is added to the address association table associated with said ingress virtual port when said address association table does not contain the source address in any destination address field of said address association table, said entry comprising the source  
20 address in a destination address field and the ingress address in a corresponding egress address field; and

on receipt of an encapsulated packet at a virtual port of the network, said encapsulated packet including a  
25 source address and an ingress address, an entry is added to the address association table associated with said virtual port when said address association table does not contain the source address in any destination address field of said address association table, said entry  
30 comprising the source address in a destination address field and the ingress address in a corresponding egress address field.

9. A method as defined in claim 1, wherein:  
35 the step of routing the packet according to the respective egress address comprises routing the packet via trunks of the network; and

when the packet is assigned a broadcast egress address corresponding to a distinct set of virtual ports, the step of routing the packet comprises routing the packet via a restricted set of trunks containing only those trunks required to reach virtual ports in the  
5 distinct set of virtual ports corresponding to said broadcast egress address.

10 10. A method as defined in claim 5, wherein:  
the step of routing the packet according to the respective egress address comprises routing the packet via trunks of the network; and

when the packet is assigned a multicast egress address corresponding to a plurality of virtual ports in a  
15 distinct set of virtual ports, the step of routing the packet comprises routing the packet via a restricted set of trunks containing only those trunks required to reach virtual ports in the plurality of virtual ports corresponding to said multicast egress address.

20 11. A communications network, comprising:  
a plurality of distinct sets of virtual ports, no virtual port belonging to more than one of the distinct sets, and each distinct set being assigned a respective  
25 distinct broadcast address;

at least one address assigner operable to assign a respective egress address to each packet entering the network via an ingress virtual port, the respective egress address corresponding to a respective destination address  
30 of the entering packet when a correspondence between the destination address and an egress address is known, and the respective egress address being a broadcast egress address corresponding to the set comprising the ingress virtual port when no correspondence between the  
35 destination address and an egress address is known; and

at least one router operable to route the packet according to the respective egress address, said routing

being restricted to virtual ports belonging to the distinct set of virtual ports which includes the ingress virtual port.

5           12.    A network as defined in claim 11, wherein, when the destination address of the packet is a unicast address and a correspondence between the destination address and a unicast egress address is known:

each address assigner is operable to assign the  
10 unicast egress address, said unicast egress address corresponding to an egress virtual port belonging to the distinct set of virtual ports which includes the ingress virtual port, the destination address being accessible from said egress virtual port; and

15           each router is operable to route the packet to said egress virtual port.

20           13.    A network as defined in claim 11, wherein, when the destination address of the packet is a unicast address and no correspondence between the destination address and an egress address is known:

each address assigner is operable to assign a  
broadcast egress address corresponding to the distinct set  
of virtual ports which includes the ingress virtual port;  
25 and

each router is operable to route the packet to  
each virtual port, other than the ingress virtual port, of  
the distinct set of virtual ports which includes the  
ingress virtual port.

30

14.    A network as defined in claim 11, wherein, when the destination address of the packet is a multicast address:

each address assigner is operable to assign a  
35 broadcast egress address corresponding to the distinct set of virtual ports which includes the ingress virtual port; and

each router is operable to route the packet to each virtual port of the distinct set of virtual ports which includes the ingress virtual port other than the ingress virtual port.

5

15. A network as defined in claim 11, wherein, when the destination address of the packet is a multicast address and a correspondence between the destination address and a multicast egress address is known:

10 each address assigner is operable to assign the multicast egress address, said multicast egress address corresponding to a plurality of virtual ports belonging to the distinct set of virtual ports which includes the ingress virtual port; and

15 each router is operable to route the packet to each virtual port of said plurality of virtual ports belonging to the distinct set of virtual ports which includes the ingress virtual port.

20 16. A network as defined in claim 11, wherein each address assigner comprises an address association table and is operable:

to assign a respective ingress address to each packet entering the network, the respective ingress address corresponding to a virtual port via which the packet enters the network;

25 to use assigned ingress addresses to populate the address association table; and

30 to use the address association table to determine correspondences between destination addresses and egress addresses.

17. A network as defined in claim 11, wherein each address assigner comprises:

35 an encapsulator for adding to each packet entering the network via an ingress virtual port the

respective egress address assigned to that packet to provide a corresponding encapsulated packet; and

a decapsulator for removing from each encapsulated packet received at an egress virtual port of the network the egress address assigned to that packet to provide a decapsulated packet.

18. A network as defined in claim 17, wherein each address assigner is operable:

to assign a respective ingress address to each packet entering the network, the respective ingress address corresponding to the ingress virtual port via which the packet enters the network;

to add the assigned ingress address to each packet entering the network in providing the corresponding encapsulated packet;

to maintain an address association table, the address association table mapping each egress address of a plurality of egress addresses to at least one corresponding destination address; and

to use the address association table to determine correspondences between destination addresses and egress addresses, wherein:

on receipt of a packet entering the network via a virtual port corresponding to an ingress address, said packet including a source address, the address assigner is operable to add an entry to the address association table when the address association table does not contain the source address in any destination address field of the address association table, said entry comprising the source address in a destination address field and the ingress address in a corresponding egress address field; and

on receipt of an encapsulated packet at a virtual port of the network, said encapsulated packet including a



source address and an ingress address, the address assigner is operable to add an entry to the address association table associated with said virtual port when said address association table does not contain the source  
5 address in any destination address field of said address association table, said entry comprising the source address in a destination address field and the ingress address in a corresponding egress address field.

10 19. A network as defined in claim 11, further comprising a plurality of trunks interconnecting routers of the network, wherein:

each router is operable to route the packet via trunks of the network; and

15 when the packet is assigned a broadcast egress address corresponding to a distinct set of virtual ports, each router is operable to route the packet via a restricted set of trunks containing only those trunks required to reach virtual ports in the distinct set of  
20 virtual ports corresponding to said broadcast egress address.

20. A network as defined in claim 15, further comprising a plurality of trunks interconnecting routers  
25 of the network, wherein:

each router is operable to route the packet via trunks of the network; and

30 when the packet is assigned a multicast egress address corresponding to a plurality of virtual ports in a distinct set of virtual ports, each router is operable to route the packet via a restricted set of trunks containing only those trunks required to reach virtual ports in the plurality of virtual ports corresponding to said multicast egress address.

35

21. A routing device for a communications network, the routing device comprising:

a plurality of distinct subsets of virtual ports,  
no virtual port belonging to more than one of the distinct  
subsets, each distinct subset being a subset of a  
respective distinct set of virtual ports of the network  
5 and each distinct set being assigned a respective distinct  
broadcast address;

at least one address assigner operable to assign  
a respective egress address to each packet entering the  
network via an ingress virtual port of the routing device,  
10 the respective egress address corresponding to a  
respective destination address of the entering packet when  
a correspondence between the destination address and an  
egress address is known, and the respective egress address  
being a broadcast egress address corresponding to the set  
15 comprising the ingress virtual port when no correspondence  
between the destination address and an egress address is  
known; and

at least one router operable to route the packet  
according to the respective egress address, said routing  
20 being restricted to virtual ports belonging to the  
distinct set of virtual ports which includes the ingress  
virtual port.

22. A routing device as defined in claim 21,  
25 wherein, when the destination address of the packet is a  
unicast address and a correspondence between the  
destination address and a unicast egress address is known:

each address assigner is operable to assign the  
unicast egress address, said unicast egress address  
30 corresponding to an egress virtual port belonging to the  
distinct set of virtual ports which includes the ingress  
virtual port, the destination address being accessible  
from said egress virtual port; and

each router is operable to route the packet to  
35 said egress virtual port.

23. A routing device as defined in claim 21, wherein, when the destination address of the packet is a unicast address and no correspondence between the destination address and an egress address is known:

5        each address assigner is operable to assign a broadcast egress address corresponding to the distinct set of virtual ports which includes the ingress virtual port; and

10        each router is operable to route the packet to each virtual port of the distinct set of virtual ports which includes the ingress virtual port other than the ingress virtual port.

24. A routing device as defined in claim 21, wherein, when the destination address of the packet is a multicast address:

15        each address assigner is operable to assign a broadcast egress address corresponding to the distinct set of virtual ports which includes the ingress virtual port; and

20        each router is operable to route the packet to each virtual port of the distinct set of virtual ports which includes the ingress virtual port other than the ingress virtual port.

25        25. A routing device as defined in claim 21, wherein, when the destination address of the packet is a multicast address and a correspondence between the destination address and a multicast egress address is known:

30        each address assigner is operable to assign the multicast egress address, said multicast egress address corresponding to a plurality of virtual ports belonging to the distinct set of virtual ports which includes the ingress virtual port; and

35        each router is operable to route the packet to each virtual port of said plurality of virtual ports

belonging to the distinct set of virtual ports which includes the ingress virtual port.

26. A routing device as defined in claim 21,  
5 wherein each address assigner comprises an address association table and is operable:

to assign a respective ingress address to each packet entering the network, the respective ingress address corresponding to a virtual port via which the  
10 packet enters the network;

to use assigned ingress addresses to populate the address association table; and

to use the address association table to determine correspondences between destination addresses and egress  
15 addresses.

27. A routing device as defined in claim 21, wherein each address assigner comprises:

an encapsulator for adding to each packet  
20 entering the network via an ingress virtual port the respective egress address assigned to that packet to provide a corresponding encapsulated packet; and

a decapsulator for removing from each encapsulated packet received at an egress virtual port of  
25 the network the egress address assigned to that packet to provide a decapsulated packet.

28. A routing device as defined in claim 27, wherein each address assigner is operable:

30 to assign a respective ingress address to each packet entering the network, the respective ingress address corresponding to the ingress virtual port via which the packet enters the network;

to add the assigned ingress address to each  
35 packet entering the network in providing the corresponding encapsulated packet;

to maintain an address association table, the address association table mapping each of a plurality of egress addresses to at least one corresponding destination address; and

5           to use the address association table to determine correspondences between destination addresses and egress addresses, wherein:

10           on receipt of a packet entering the network via a virtual port associated with an ingress address, said packet including a source address, the address assigner is operable to add an entry to the address association table when the address association table does not contain the source address in any destination address field of the  
15           address association table, said entry comprising the source address in a destination address field and the ingress address in a corresponding egress address field; and

20           on receipt of an encapsulated packet via a virtual port of the network, said encapsulated packet including a source address and an ingress address, the address assigner is operable to add an entry to the address association table associated with said virtual  
25           port when said address association table does not contain the source address in any destination address field of said address association table, said entry comprising the source address in a destination address field and the ingress address in a corresponding egress address field.

30

29. A routing device as defined in claim 21, wherein:

each router is operable to route the packet via trunks of the network; and

35           when the packet is assigned a broadcast egress address corresponding to a distinct set of virtual ports, each router is operable to route the packet via a

restricted set of trunks containing only those trunks required to reach virtual ports in the distinct set of virtual ports corresponding to said broadcast egress address.

5

30. A routing device as defined in claim 25, wherein:

each router is operable to route the packet via trunks of the network; and

10 when the packet is assigned a multicast egress address corresponding to a plurality of virtual ports in a distinct set of virtual ports, each router is operable to route the packet via a restricted set of trunks containing only those trunks required to reach virtual ports in the  
15 plurality of virtual ports corresponding to said multicast egress address.

31. A routing device as defined in claim 28, wherein each router provides IEEE 802.1 switching  
20 functionality adapted to packets encapsulated with ingress and egress addresses.

32. A routing device as defined in claim 28, comprising a respective address assigner for each distinct  
25 subset of virtual ports, each address assigner being connected between its respective distinct subset of virtual ports and a router of the routing device.

33. A routing device as defined in claim 32,  
30 further comprising a switching element connected between at least one address assigner and its respective distinct subset of virtual ports, said switching element being operable to multiplex the virtual ports of the respective distinct subset of virtual ports onto the address  
35 assigner.

34. A routing device as defined in claim 33,  
wherein:

each switching element provides IEEE 802.1  
switching functionality; and

5 each router provides IEEE 802.1 switching  
functionality adapted to packets encapsulated with ingress  
and egress addresses.

35. A routing device as defined in claim 32,  
10 further comprising a plurality of VLAN demultiplexers  
connected to the router, each VLAN demultiplexer being  
connected between the router and a respective plurality of  
the address assigners, each VLAN demultiplexer being  
associated with a respective egress address and being  
15 operable to route an encapsulated packet from the router  
to an address assigner associated with the ingress address  
of the encapsulated packet such that all encapsulated  
packets having a common egress address and an ingress  
address corresponding to a virtual port in a particular  
20 set of the distinct sets of virtual ports are routed to an  
address assigner associated with that egress address and  
that particular distinct set of virtual ports.

36. A routing device as defined in claim 35,  
25 further comprising:

a respective VLAN translator connected to each  
address assigner that is connected to the VLAN  
demultiplexer, each VLAN translator being operable to  
apply a respective VLAN identifier to packets received  
30 from its respective address assigner; and

a router demultiplexer connected to a plurality  
of the VLAN translators for routing packets received from  
an external router to a VLAN translator selected according  
to VLAN identifiers of the packets received from the  
35 external router.

37. A routing device as defined in claim 35, further comprising a respective virtual private router connected to each address assigner that is connected to a VLAN demultiplexer.

5

38. A routing device as defined in claim 37, further comprising a respective network address translator connected to each virtual private router for translating addresses between a respective first address space used by  
10 its virtual private router and a second address space used by an Internet router.

39. A routing device as defined in claim 38, further comprising an Internet router connected to the  
15 network address translators.

40. A routing device as defined in claim 35, further comprising:

an MPLS switch, the MPLS switch being operable to  
20 route packets between an Internet router and address assigners selected according to MPLS labels of the packets; and

a respective MPLS converter connected between each address assigner that is connected to a VLAN  
25 demultiplexer and the MPLS switch, each MPLS converter:

being operable to apply a respective MPLS label to each packet received from its respective address assigner, said MPLS label being uniquely associated with the MPLS converter; and

30 being operable to remove MPLS labels from packets received from the MPLS switch.

41. A method as defined in claim 8, further comprising routing an encapsulated packet from the router  
35 to an address assigner selected according to the ingress address and the egress address of the encapsulated packet such that all encapsulated packets having a common egress



address and an ingress address corresponding to a virtual port in a particular set of the distinct sets of virtual ports are routed to an address assigner associated with that egress address and that particular distinct set of  
5 virtual ports.

42. A method as defined in claim 41, further comprising:

applying a respective VLAN identifier to packets  
10 leaving the network from a respective address assigner;  
and

routing packets received from an external router to an address assigner selected according to VLAN identifiers of the packets received from the external  
15 router.

43. A method as defined in claim 41, further comprising:

applying a respective MPLS label to packets  
20 leaving the network from an address assigner, said MPLS label being uniquely associated with said address assigner;

routing packets between an Internet router and address assigners according to MPLS labels of the packets;  
25 and

removing MPLS labels from packets received from the Internet router.

44. A method as defined in claim 41, further comprising:  
30

applying a respective identifier to packets leaving the network from an address assigner, said identifier being uniquely associated with said address assigner; and

35 routing packets into and out of the network according to their respective identifiers.

45. A method as defined in claim 1, wherein at least one physical port of the network maps one-to-one onto a corresponding virtual port of network, said physical port and said corresponding virtual port being  
5 associated with a respective distinct physical address.

46. A method as defined in claim 1, wherein at least one physical port of the network maps onto a corresponding plurality of virtual ports of the network,  
10 said physical port being associated with a respective distinct physical address, and each virtual port of said corresponding plurality of virtual ports being associated with a respective distinct combination of said physical address and a respective virtual network identifier.

47. A network as defined in claim 11, wherein at least one physical port of the network maps one-to-one onto a corresponding virtual port of network, said physical port and said corresponding virtual port being  
20 associated with a respective distinct physical address.

48. A network as defined in claim 11, wherein at least one physical port of the network maps onto a corresponding plurality of virtual ports of the network,  
25 said physical port being associated with a respective distinct physical address, and each virtual port of said corresponding plurality of virtual ports being associated with a respective distinct combination of said physical address and a respective virtual network identifier.

49. A routing device as defined in claim 21, wherein at least one physical port of the routing device maps one-to-one onto a corresponding virtual port of routing device, said physical port and said corresponding  
35 virtual port being associated with a respective distinct physical address.

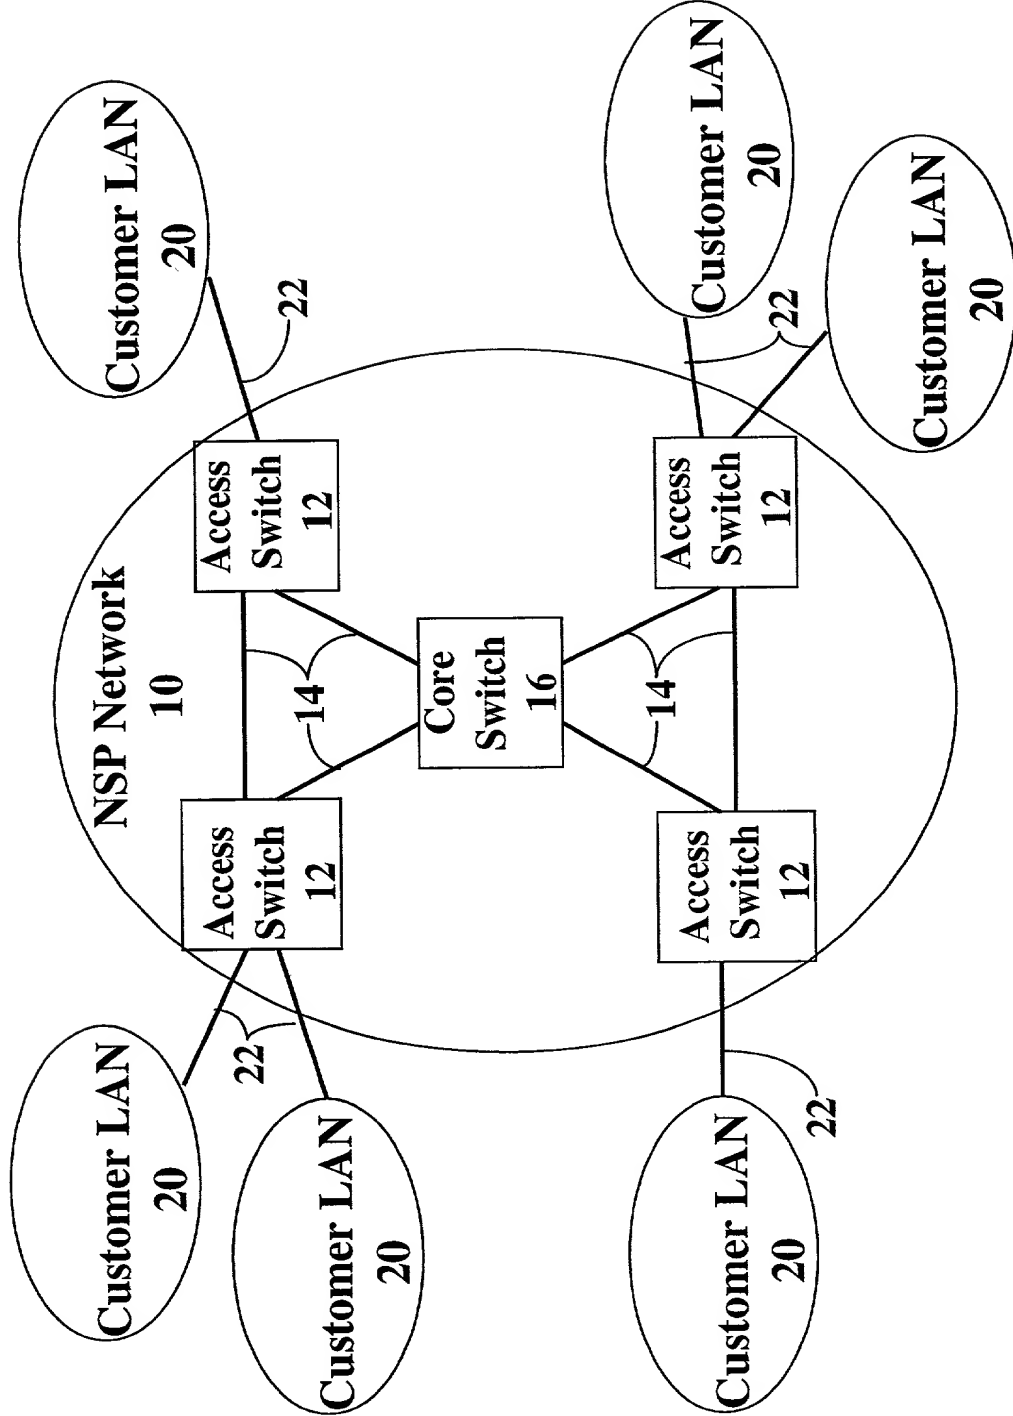
50. A routing device as defined in claim 21,  
wherein at least one physical port of the routing device  
maps onto a corresponding plurality of virtual ports of  
the routing device, said physical port being associated  
5 with a respective distinct physical address, and each  
virtual port of said corresponding plurality of virtual  
ports being associated with a respective distinct  
combination of said physical address and a respective  
virtual network identifier.

10

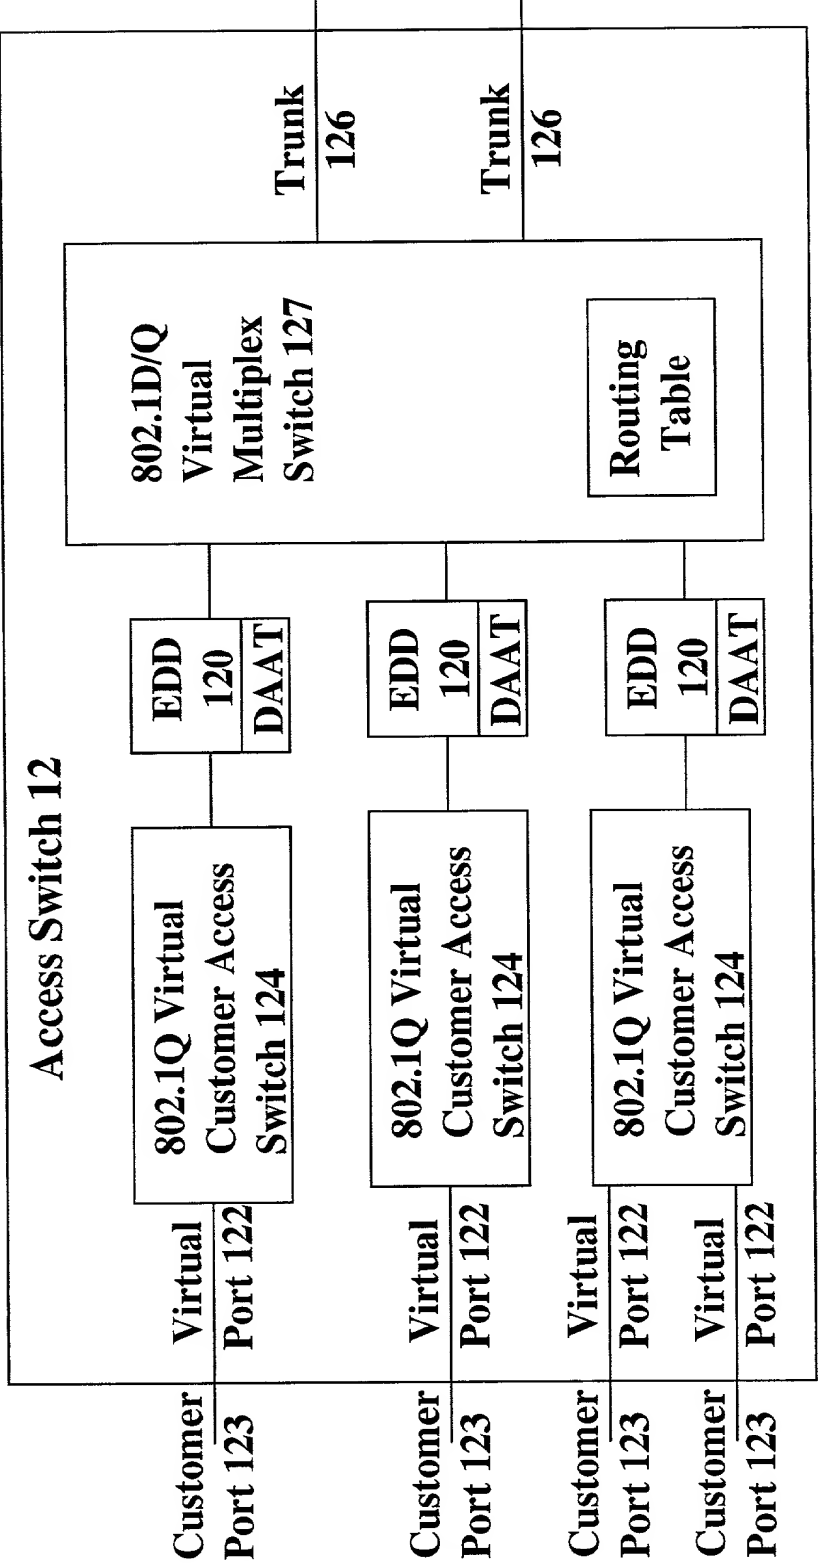
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65  
66  
67  
68  
69  
70  
71  
72  
73  
74  
75  
76  
77  
78  
79  
80  
81  
82  
83  
84  
85  
86  
87  
88  
89  
90  
91  
92  
93  
94  
95  
96  
97  
98  
99  
100  
101  
102  
103  
104  
105  
106  
107  
108  
109  
110  
111  
112  
113  
114  
115  
116  
117  
118  
119  
120  
121  
122  
123  
124  
125  
126  
127  
128  
129  
130  
131  
132  
133  
134  
135  
136  
137  
138  
139  
140  
141  
142  
143  
144  
145  
146  
147  
148  
149  
150  
151  
152  
153  
154  
155  
156  
157  
158  
159  
160  
161  
162  
163  
164  
165  
166  
167  
168  
169  
170  
171  
172  
173  
174  
175  
176  
177  
178  
179  
180  
181  
182  
183  
184  
185  
186  
187  
188  
189  
190  
191  
192  
193  
194  
195  
196  
197  
198  
199  
200  
201  
202  
203  
204  
205  
206  
207  
208  
209  
210  
211  
212  
213  
214  
215  
216  
217  
218  
219  
220  
221  
222  
223  
224  
225  
226  
227  
228  
229  
230  
231  
232  
233  
234  
235  
236  
237  
238  
239  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
250  
251  
252  
253  
254  
255  
256  
257  
258  
259  
260  
261  
262  
263  
264  
265  
266  
267  
268  
269  
270  
271  
272  
273  
274  
275  
276  
277  
278  
279  
280  
281  
282  
283  
284  
285  
286  
287  
288  
289  
290  
291  
292  
293  
294  
295  
296  
297  
298  
299  
300  
301  
302  
303  
304  
305  
306  
307  
308  
309  
310  
311  
312  
313  
314  
315  
316  
317  
318  
319  
320  
321  
322  
323  
324  
325  
326  
327  
328  
329  
330  
331  
332  
333  
334  
335  
336  
337  
338  
339  
340  
341  
342  
343  
344  
345  
346  
347  
348  
349  
350  
351  
352  
353  
354  
355  
356  
357  
358  
359  
360  
361  
362  
363  
364  
365  
366  
367  
368  
369  
370  
371  
372  
373  
374  
375  
376  
377  
378  
379  
380  
381  
382  
383  
384  
385  
386  
387  
388  
389  
390  
391  
392  
393  
394  
395  
396  
397  
398  
399  
400  
401  
402  
403  
404  
405  
406  
407  
408  
409  
410  
411  
412  
413  
414  
415  
416  
417  
418  
419  
420  
421  
422  
423  
424  
425  
426  
427  
428  
429  
430  
431  
432  
433  
434  
435  
436  
437  
438  
439  
440  
441  
442  
443  
444  
445  
446  
447  
448  
449  
450  
451  
452  
453  
454  
455  
456  
457  
458  
459  
460  
461  
462  
463  
464  
465  
466  
467  
468  
469  
470  
471  
472  
473  
474  
475  
476  
477  
478  
479  
480  
481  
482  
483  
484  
485  
486  
487  
488  
489  
490  
491  
492  
493  
494  
495  
496  
497  
498  
499  
500  
501  
502  
503  
504  
505  
506  
507  
508  
509  
510  
511  
512  
513  
514  
515  
516  
517  
518  
519  
520  
521  
522  
523  
524  
525  
526  
527  
528  
529  
530  
531  
532  
533  
534  
535  
536  
537  
538  
539  
540  
541  
542  
543  
544  
545  
546  
547  
548  
549  
550  
551  
552  
553  
554  
555  
556  
557  
558  
559  
560  
561  
562  
563  
564  
565  
566  
567  
568  
569  
570  
571  
572  
573  
574  
575  
576  
577  
578  
579  
580  
581  
582  
583  
584  
585  
586  
587  
588  
589  
590  
591  
592  
593  
594  
595  
596  
597  
598  
599  
600  
601  
602  
603  
604  
605  
606  
607  
608  
609  
610  
611  
612  
613  
614  
615  
616  
617  
618  
619  
620  
621  
622  
623  
624  
625  
626  
627  
628  
629  
630  
631  
632  
633  
634  
635  
636  
637  
638  
639  
640  
641  
642  
643  
644  
645  
646  
647  
648  
649  
650  
651  
652  
653  
654  
655  
656  
657  
658  
659  
660  
661  
662  
663  
664  
665  
666  
667  
668  
669  
670  
671  
672  
673  
674  
675  
676  
677  
678  
679  
680  
681  
682  
683  
684  
685  
686  
687  
688  
689  
690  
691  
692  
693  
694  
695  
696  
697  
698  
699  
700  
701  
702  
703  
704  
705  
706  
707  
708  
709  
710  
711  
712  
713  
714  
715  
716  
717  
718  
719  
720  
721  
722  
723  
724  
725  
726  
727  
728  
729  
730  
731  
732  
733  
734  
735  
736  
737  
738  
739  
740  
741  
742  
743  
744  
745  
746  
747  
748  
749  
750  
751  
752  
753  
754  
755  
756  
757  
758  
759  
760  
761  
762  
763  
764  
765  
766  
767  
768  
769  
770  
771  
772  
773  
774  
775  
776  
777  
778  
779  
780  
781  
782  
783  
784  
785  
786  
787  
788  
789  
790  
791  
792  
793  
794  
795  
796  
797  
798  
799  
800  
801  
802  
803  
804  
805  
806  
807  
808  
809  
810  
811  
812  
813  
814  
815  
816  
817  
818  
819  
820  
821  
822  
823  
824  
825  
826  
827  
828  
829  
830  
831  
832  
833  
834  
835  
836  
837  
838  
839  
840  
841  
842  
843  
844  
845  
846  
847  
848  
849  
850  
851  
852  
853  
854  
855  
856  
857  
858  
859  
860  
861  
862  
863  
864  
865  
866  
867  
868  
869  
870  
871  
872  
873  
874  
875  
876  
877  
878  
879  
880  
881  
882  
883  
884  
885  
886  
887  
888  
889  
890  
891  
892  
893  
894  
895  
896  
897  
898  
899  
900  
901  
902  
903  
904  
905  
906  
907  
908  
909  
910  
911  
912  
913  
914  
915  
916  
917  
918  
919  
920  
921  
922  
923  
924  
925  
926  
927  
928  
929  
930  
931  
932  
933  
934  
935  
936  
937  
938  
939  
940  
941  
942  
943  
944  
945  
946  
947  
948  
949  
950  
951  
952  
953  
954  
955  
956  
957  
958  
959  
960  
961  
962  
963  
964  
965  
966  
967  
968  
969  
970  
971  
972  
973  
974  
975  
976  
977  
978  
979  
980  
981  
982  
983  
984  
985  
986  
987  
988  
989  
990  
991  
992  
993  
994  
995  
996  
997  
998  
999  
1000  
1001  
1002  
1003  
1004  
1005  
1006  
1007  
1008  
1009  
1010  
1011  
1012  
1013  
1014  
1015  
1016  
1017  
1018  
1019  
1020  
1021  
1022  
1023  
1024  
1025  
1026  
1027  
1028  
1029  
1030  
1031  
1032  
1033  
1034  
1035  
1036  
1037  
1038  
1039  
1040  
1041  
1042  
1043  
1044  
1045  
1046  
1047  
1048  
1049  
1050  
1051  
1052  
1053  
1054  
1055  
1056  
1057  
1058  
1059  
1060  
1061  
1062  
1063  
1064  
1065  
1066  
1067  
1068  
1069  
1070  
1071  
1072  
1073  
1074  
1075  
1076  
1077  
1078  
1079  
1080  
1081  
1082  
1083  
1084  
1085  
1086  
1087  
1088  
1089  
1090  
1091  
1092  
1093  
1094  
1095  
1096  
1097  
1098  
1099  
1100  
1101  
1102  
1103  
1104  
1105  
1106  
1107  
1108  
1109  
1110  
1111  
1112  
1113  
1114  
1115  
1116  
1117  
1118  
1119  
1120  
1121  
1122  
1123  
1124  
1125  
1126  
1127  
1128  
1129  
1130  
1131  
1132  
1133  
1134  
1135  
1136  
1137  
1138  
1139  
1140  
1141  
1142  
1143  
1144  
1145  
1146  
1147  
1148  
1149  
1150  
1151  
1152  
1153  
1154  
1155  
1156  
1157  
1158  
1159  
1160  
1161  
1162  
1163  
1164  
1165  
1166  
1167  
1168  
1169  
1170  
1171  
1172  
1173  
1174  
1175  
1176  
1177  
1178  
1179  
1180  
1181  
1182  
1183  
1184  
1185  
1186  
1187  
1188  
1189  
1190  
1191  
1192  
1193  
1194  
1195  
1196  
1197  
1198  
1199  
1200  
1201  
1202  
1203  
1204  
1205  
1206  
1207  
1208  
1209  
1210  
1211  
1212  
1213  
1214  
1215  
1216  
1217  
1218  
1219  
1220  
1221  
1222  
1223  
1224  
1225  
1226  
1227  
1228  
1229  
1230  
1231  
1232  
1233  
1234  
1235  
1236  
1237  
1238  
1239  
1240  
1241  
1242  
1243  
1244  
1245  
1246  
1247  
1248  
1249  
1250  
1251  
1252  
1253  
1254  
1255  
1256  
1257  
1258  
1259  
1260  
1261  
1262  
1263  
1264  
1265  
1266  
1267  
1268  
1269  
1270  
1271  
1272  
1273  
1274  
1275  
1276  
1277  
1278  
1279  
1280  
1281  
1282  
1283  
1284  
1285  
1286  
1287  
1288  
1289  
1290  
1291  
1292  
1293  
1294  
1295  
1296  
1297  
1298  
1299  
1300  
1301  
1302  
1303  
1304  
1305  
1306  
1307  
1308  
1309  
1310  
1311  
1312  
1313  
1314  
1315  
1316  
1317  
1318  
1319  
1320  
1321  
1322  
1323  
1324  
1325  
1326  
1327  
1328  
1329  
1330  
1331  
1332  
1333  
1334  
1335  
1336  
1337  
1338  
1339  
1340  
1341  
1342  
1343  
1344  
1345  
1346  
1347  
1348  
1349  
1350  
1351  
1352  
1353  
1354  
1355  
1356  
1357  
1358  
1359  
1360  
1361  
1362  
1363  
1364  
1365  
1366  
1367  
1368  
1369  
1370  
1371  
1372  
1373  
1374  
1375  
1376  
1377  
1378  
1379  
1380  
1381  
1382  
1383  
1384  
1385  
1386  
1387  
1388  
1389  
1390  
1391  
1392  
1393  
1394  
1395  
1396  
1397  
1398  
1399  
1400  
1401  
1402  
1403  
1404  
1405  
1406  
1407  
1408  
1409  
1410  
1411  
1412  
1413  
1414  
1415  
1416  
1417  
1418  
1419  
1420  
1421  
1422  
1423  
1424  
1425  
1426  
1427  
1428  
1429  
1430  
1431  
1432  
1433  
1434  
1435  
1436  
1437  
1438  
1439  
1440  
1441  
1442  
1443  
1444  
1445  
1446  
1447  
1448  
1449  
1450  
1451  
1452  
1453  
1454  
1455  
1456  
1457  
1458  
1459  
1460  
1461  
1462  
1463  
1464  
1465  
1466  
1467  
1468  
1469  
1470  
1471  
1472  
1473  
1474  
1475  
1476  
1477  
1478  
1479  
1480  
1481  
1482  
1483  
1484  
1485  
1486  
1487  
1488  
1489  
1490  
1491  
1492  
1493  
1494  
1495  
1496  
1497  
1498  
1499  
1500  
1501  
1502  
1503  
1504  
1505  
1506  
1507  
1508  
1509  
1510  
1511  
1512  
1513  
1514  
1515  
1516  
1517  
1518  
1519  
1520  
1521  
1522  
1523  
1524  
1525  
1526  
1527  
1528  
1529  
1530  
1531  
1532  
1533  
1534  
1535  
1536  
1537  
1538  
1539  
1540  
1541  
1542  
1543  
1544  
1545  
1546  
1547  
1548  
1549  
1550  
1551  
1552  
1553  
1554  
1555  
1556  
1557  
1558  
1559  
1560  
1561  
1562  
1563  
1564  
1565  
1566  
1567  
1568  
1569  
1570  
1571  
1572  
1573  
1574  
1575  
1576  
1577  
1578  
1579  
1580  
1581  
1582  
1583  
1584  
1585  
1586  
1587  
1588  
1589  
1590  
1591  
1592  
1593  
1594  
1595  
1596  
1597  
1598  
1599  
1600  
1601  
1602  
1603  
1604  
1605  
1606  
1607  
1608  
1609  
1610  
1611  
1612  
1613  
1614  
1615  
1616  
1617  
1618  
1619  
1620  
1621  
1622  
1623  
1624  
1625  
1626  
1627  
1628  
1629  
1630  
1631  
1632  
1633  
1634  
1635  
1636  
1637  
1638  
1639  
1640  
1641  
1642  
1643  
1644  
1645  
1646  
1647  
1648  
1649  
1650  
1651  
1652  
1653  
1654  
1655  
1656  
1657  
1658  
1659  
1660  
1661  
1662  
1663  
1664  
1665  
1666  
1667  
1668  
1669  
1670  
1671  
1672  
1673  
1674  
1675  
1676  
1677  
1678  
1679  
1680  
1681  
1682  
1683  
1684  
1685  
1686  
1687  
1688  
1689  
1690  
1691  
1692  
1693  
1694  
1695  
1696  
1697  
1698  
1699  
1700  
1701  
1702  
1703  
1704  
1705  
1706  
1707  
1708  
1709  
1710  
1711  
1712  
1713  
1714  
1715  
1716  
1717  
1718  
1719  
1720  
1721  
1722  
1723  
1724  
1725  
1726  
1727  
1728  
1729  
1730  
1731  
1732  
1733  
1734  
1735  
1736  
1737  
1738  
1739  
1740  
1741  
1742  
1743  
1744  
1745  
1746  
1747  
1748  
1749  
1750  
1751  
1752  
1753  
1754  
1755  
1756  
1757  
1758  
1759  
1760  
1761  
1762  
1763  
1764  
1765  
1766  
1767  
1768  
1769  
1770  
1771  
1772  
1773  
1774  
1775  
1776  
1777  
1778  
1779  
1780  
1781  
1782  
1783  
1784  
1785  
1786  
1787  
1788  
1789  
1790  
1791  
1792  
1793  
1794  
1795  
1796  
1797  
1798  
1799  
1800  
1801  
1802  
1803  
1804  
1805  
1806  
1807  
1808  
1809  
1810  
1811  
1812  
1813  
1814  
1815  
1816  
1817  
1818  
1819  
1820  
1821  
1822  
1823  
1824  
1825  
1826  
1827  
1828  
1829  
1830  
1831  
1832  
1833  
1834  
1835  
1836  
1837  
1838  
1839  
1840  
1841  
1842  
1843  
1844  
1845  
1846  
1847  
1848  
1849  
1850  
1851  
1852  
1853  
1854  
1855  
1856  
1857  
1858  
1859  
1860  
1861  
1862  
1863  
1864  
1865  
1866  
1867  
1868  
1869  
1870  
1871  
1872  
1873  
1874  
1875  
1876  
1877  
1878  
1879  
1880  
1881  
1882  
1883  
1884  
1885  
1886  
1887  
1888  
1889  
1890  
1891  
1892  
1893  
1894  
1895  
1896  
1897  
1898  
1899  
1900  
1901  
1902  
1903  
1904  
1905  
1906  
1907  
1908  
1909  
1910  
1911  
1912  
1913  
1914  
1915  
1916  
1917  
1918  
1919  
1920  
1921  
1922  
1923  
1924  
1925  
1926  
1927  
1928  
1929  
1930  
1931  
1932  
1933  
1934  
1935  
1936  
1937  
1938  
1939  
1940  
1941  
1942  
1943  
1944  
1945  
1946  
1947  
1948  
1949  
1950  
1951  
1952  
1953  
1954  
1955  
1956  
1957  
1958  
1959  
1960  
1961  
1962  
1963  
1964  
1965  
1966  
1967  
1968  
1969  
1970  
1971  
1972  
1973  
1974  
1975  
1976  
1977  
1978  
1979  
1980  
1981  
1982  
1983  
1984  
1985  
1986  
1987  
1988  
1989  
1990  
1991  
1992  
1993  
1994  
1995  
1996  
1997  
1998  
1999  
2000  
2001  
2002  
2003  
2004  
2005  
2006  
2007  
2008  
2009  
2010  
2011  
2012  
2013  
2014  
2015  
2016  
2017  
2018  
2019  
2020  
2021  
2022  
2023  
2024  
2025  
2026  
2027  
2028  
2029  
2030  
2031  
2032  
2033  
2034  
2035  
2036  
2037  
2038  
2039  
2040  
2041  
2042  
2043  
2044  
2045  
2046  
2047  
2048  
2049  
2050  
2051  
2052  
2053  
2054  
2055  
2056  
2057  
2058  
2059  
2060  
2061  
2062  
2063  
2064  
2065  
2066  
2067  
2068  
2069  
2070  
2071  
2072  
2073  
2074  
2075  
2076  
2077  
2078  
2079  
2080  
2081  
2082  
2083  
2084  
2085  
2086  
2087  
2088  
2089  
2090  
2091  
2092  
2093  
2094  
2095  
2096  
2097  
2098  
2099  
2100  
2101  
2102  
2103  
2104  
2105  
2106  
2107  
2108  
2109  
2110  
2111  
2112  
2113  
2114  
2115  
2116  
2117  
2118  
2119  
2120  
2121  
2122  
2123  
2124  
2125  
2126  
2127  
2128  
2129  
2130  
2131  
2132  
2133  
2134  
2135  
2136  
2137  
2138  
2139  
2140  
2141  
2142  
2143  
2144  
2145  
2146  
2147  
2148  
2149  
2150  
2151  
2152  
2153  
2154  
2155  
2156  
2157  
2158  
2159  
2160  
2161  
2162  
2163  
2164  
2165  
2166  
2167  
2168  
2169  
2170  
2171  
2172  
2173  
2174  
2175  
2176  
2177  
2178  
2179  
2180  
2181  
2182  
2183  
2184  
2185  
2186  
2187  
2188  
2189  
2190  
2191  
2192  
2193  
2194  
2195  
2196  
2197  
2198  
2199  
2200  
2201  
2202  
2203  
2204  
2205  
2206  
2207  
2208  
2209  
2210  
2211  
2212  
2213  
2214  
2215  
2216

Abstract of the Disclosure

In methods and apparatus for routing packets through a communications network, a respective distinct broadcast address is assigned to each of a plurality of  
5 distinct sets of virtual ports. No virtual port belongs to more than one of the distinct sets. A respective egress address is assigned to each packet entering the network via an ingress virtual port. The respective egress address corresponds to a respective destination address of  
10 the entering packet when a correspondence between the destination address and an egress address is known. When no correspondence between the destination address and an egress address is known, the respective egress address is a broadcast egress address corresponding to the set  
15 comprising the ingress virtual port. The packet is routed according to the respective egress address. The routing is restricted to virtual ports belonging to the distinct set of virtual ports that includes the ingress virtual port. The distinct sets of virtual ports and their  
20 associated broadcast addresses define isolated virtual private networks within the network. Each physical port of the network may map one-to-one onto a corresponding virtual port, or may map onto a corresponding plurality of virtual ports, in which case the each virtual port of the  
25 plurality is associated with a respective distinct combination of a physical address of the physical port and a respective virtual network identifier.



**Fig. 1**



**Fig. 2**

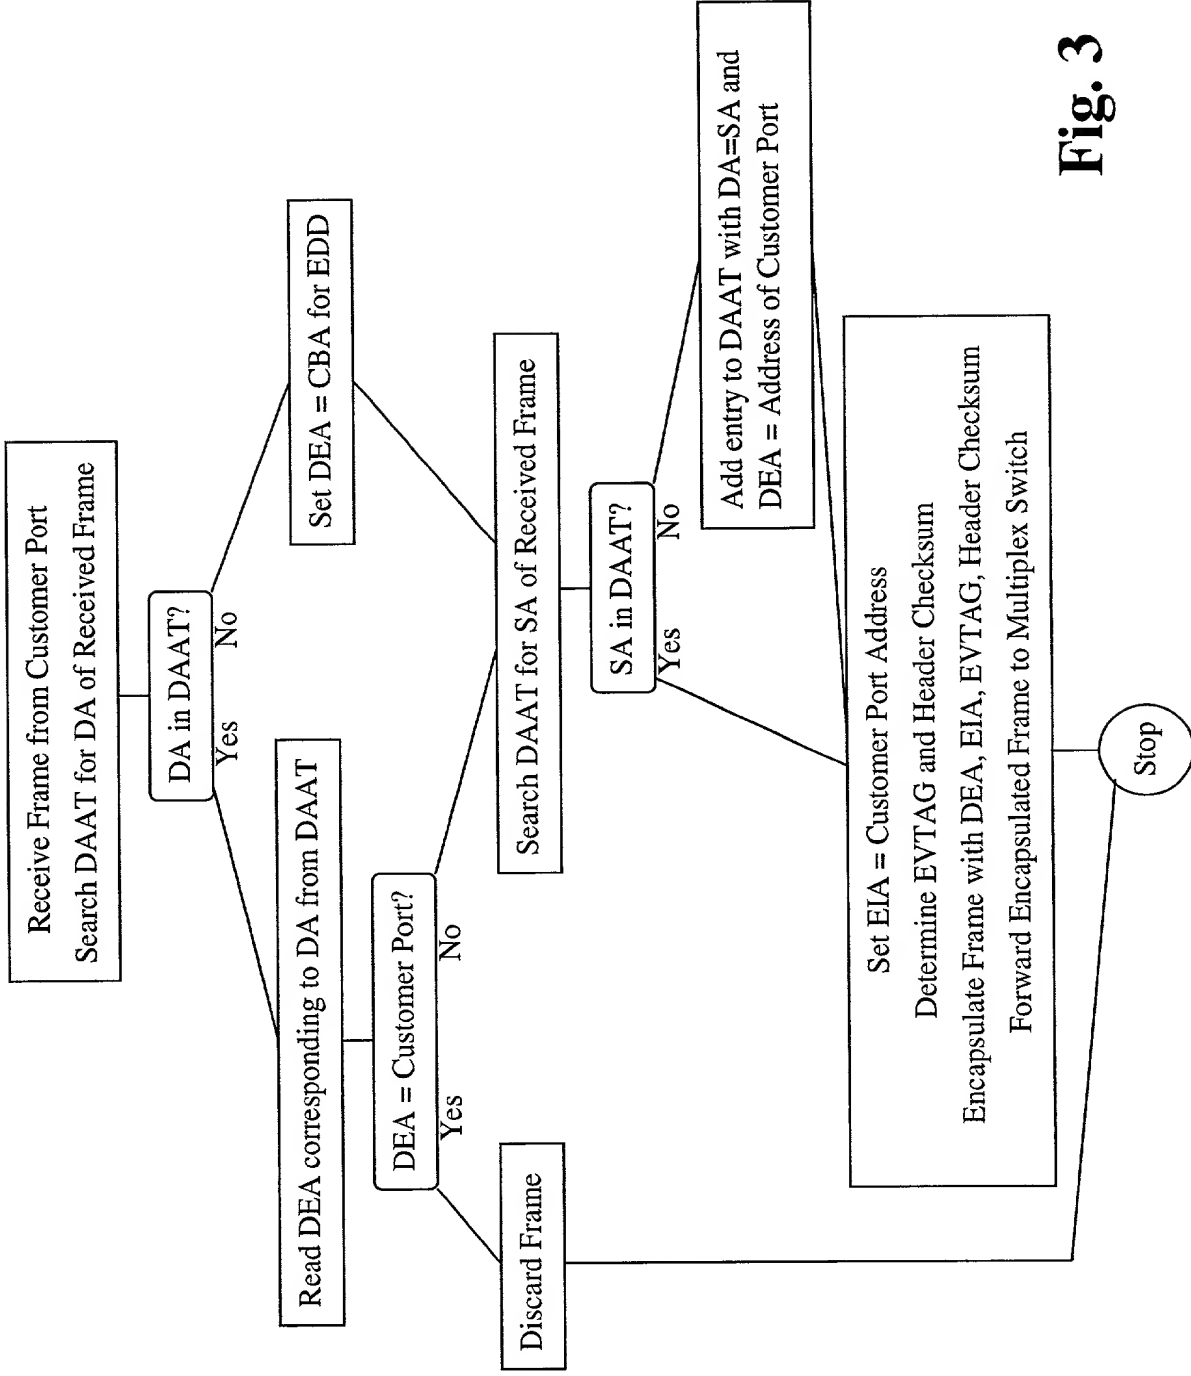
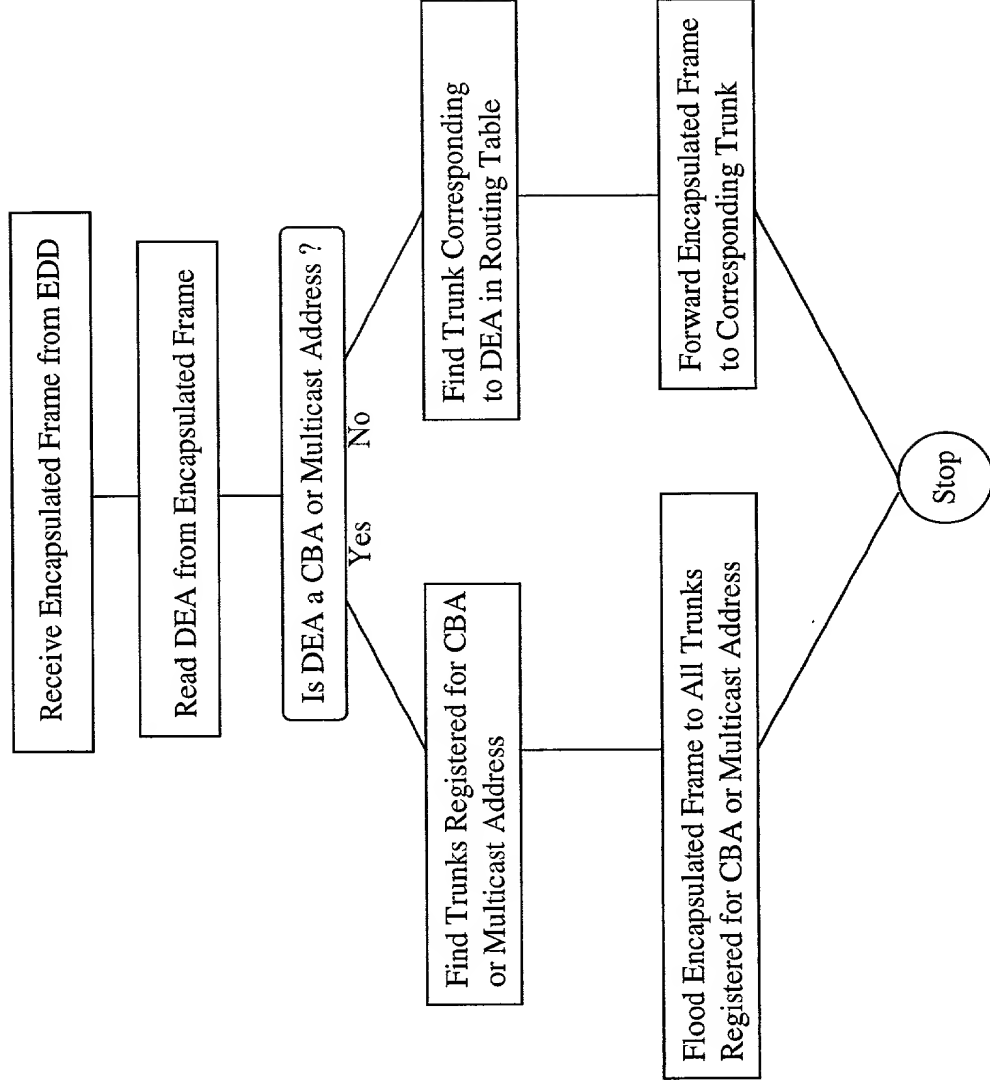
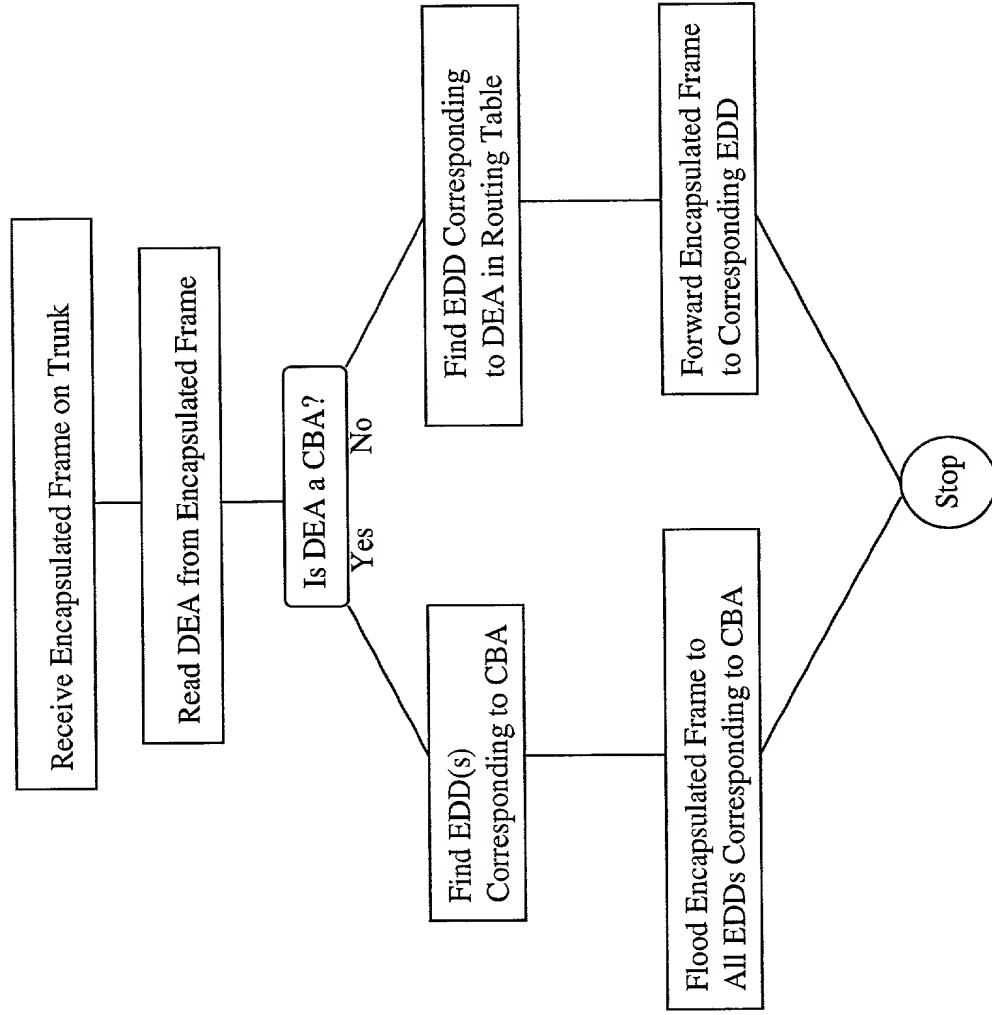


Fig. 3

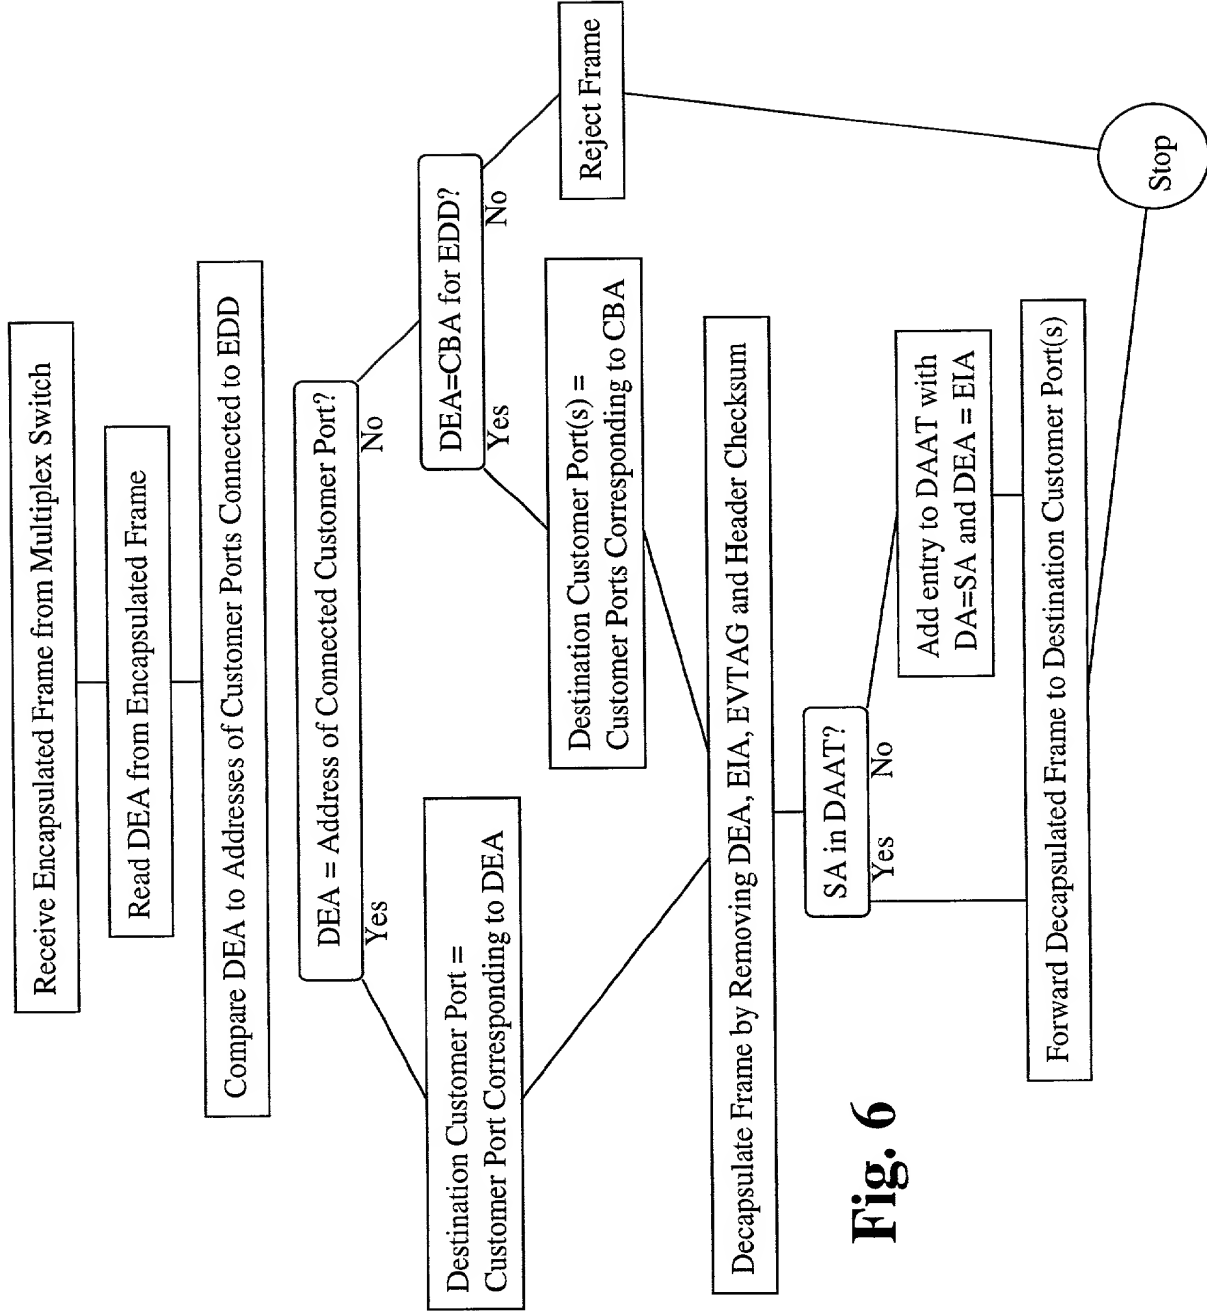


**Fig. 4**





**Fig. 5**



**Fig. 6**


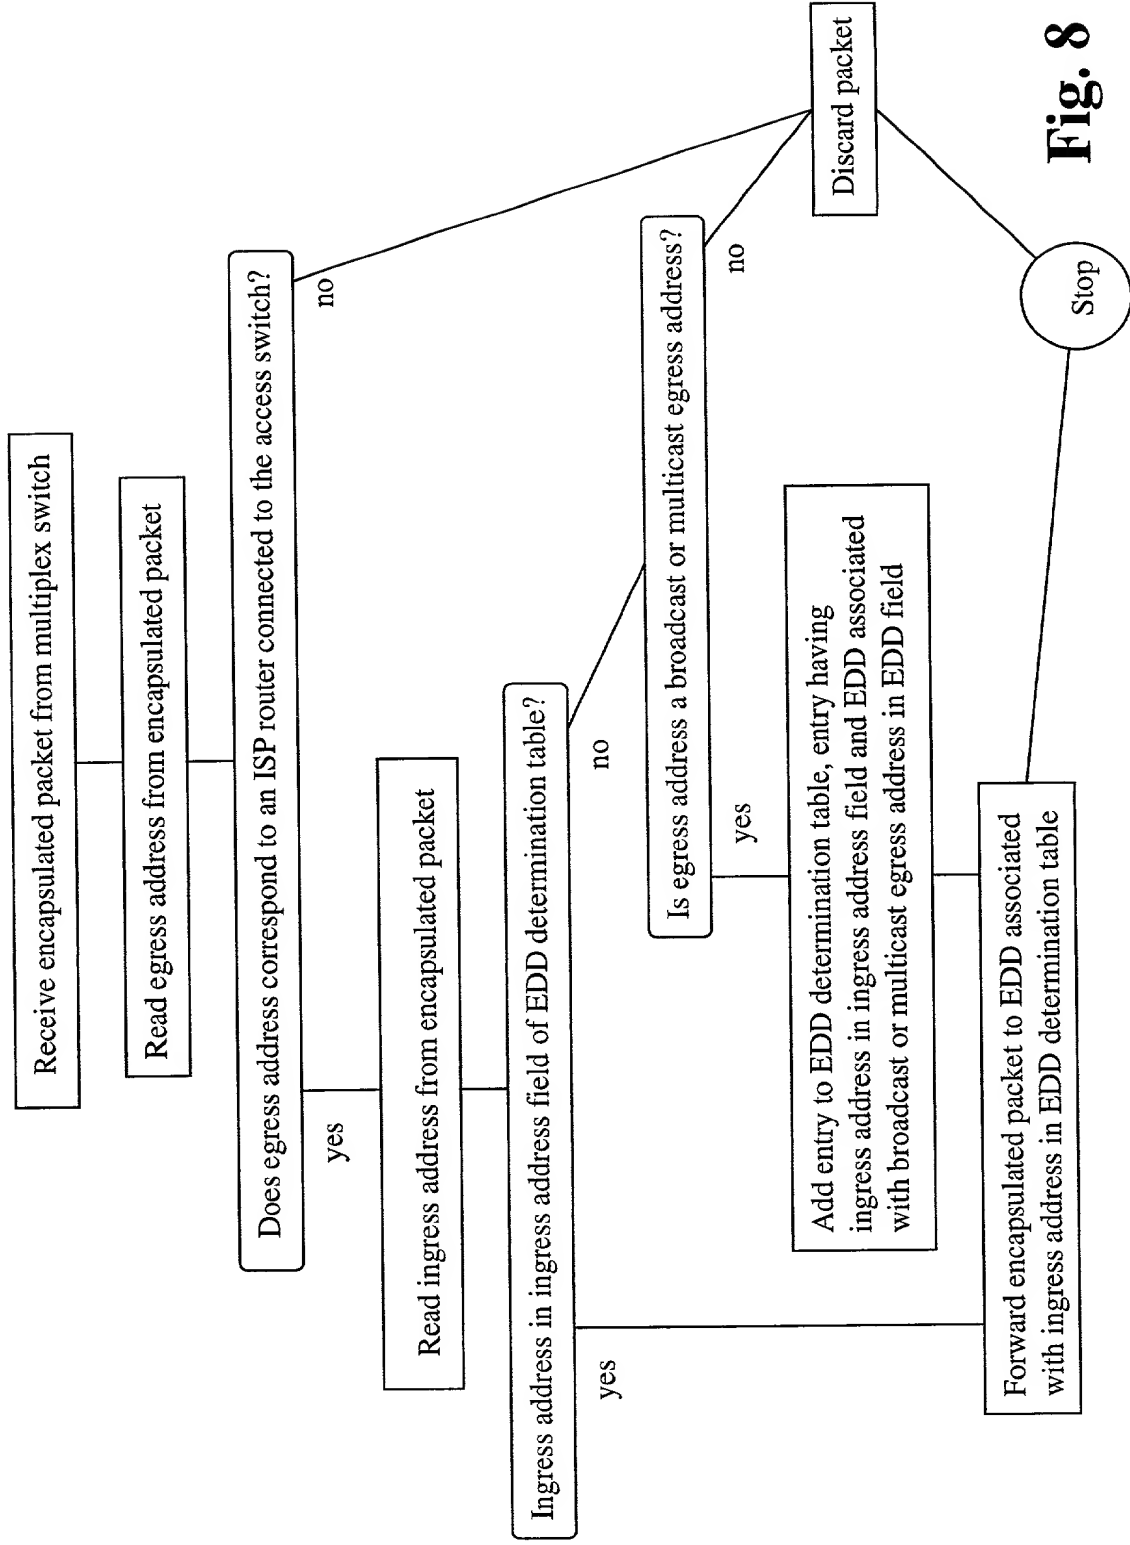


Fig. 7



**Fig. 8**

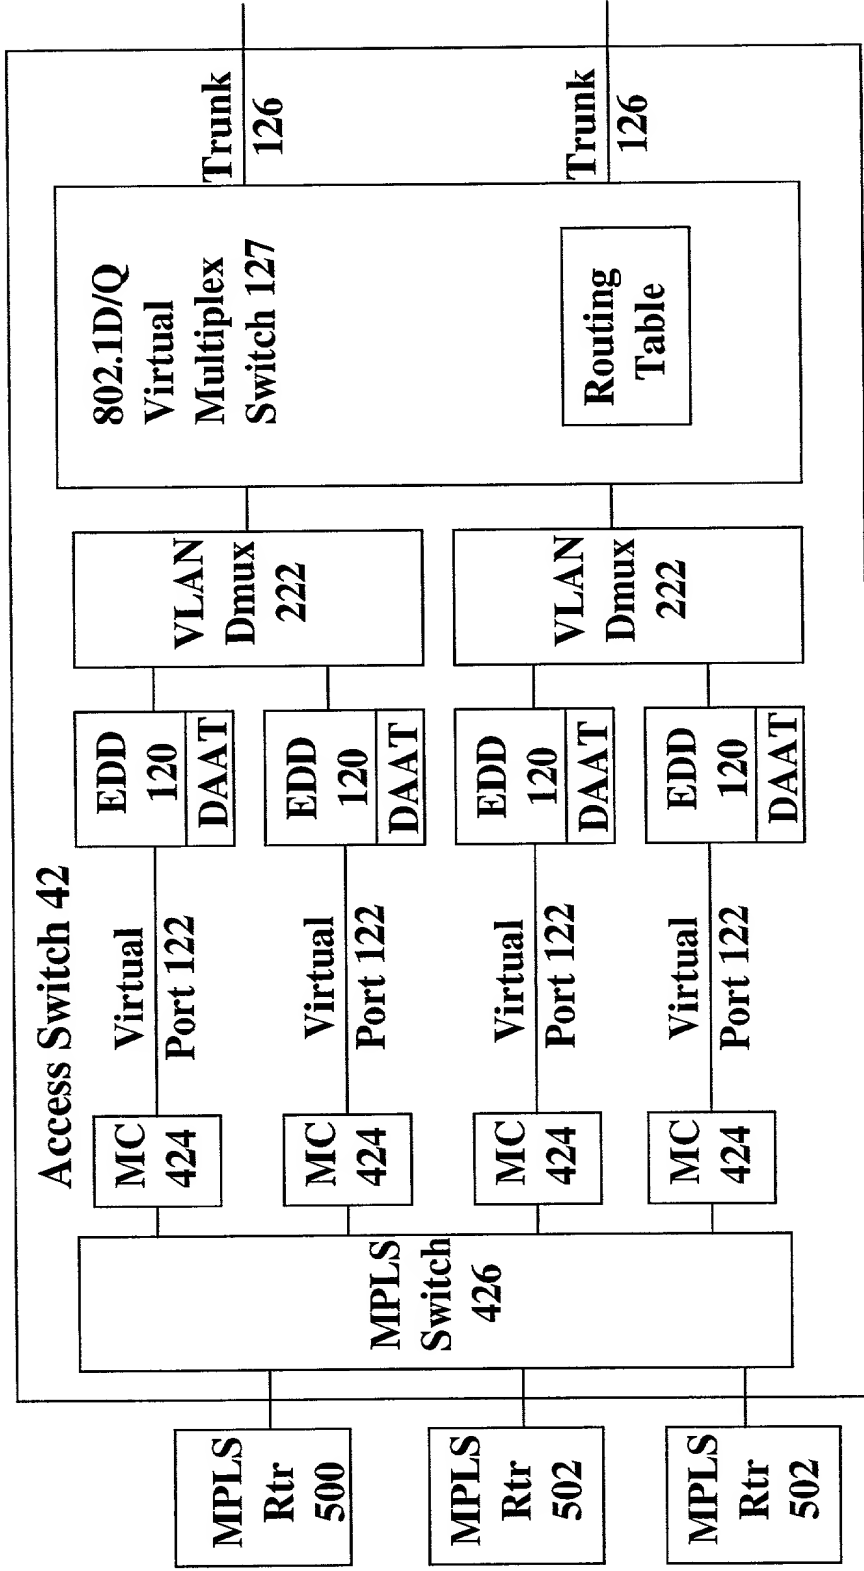
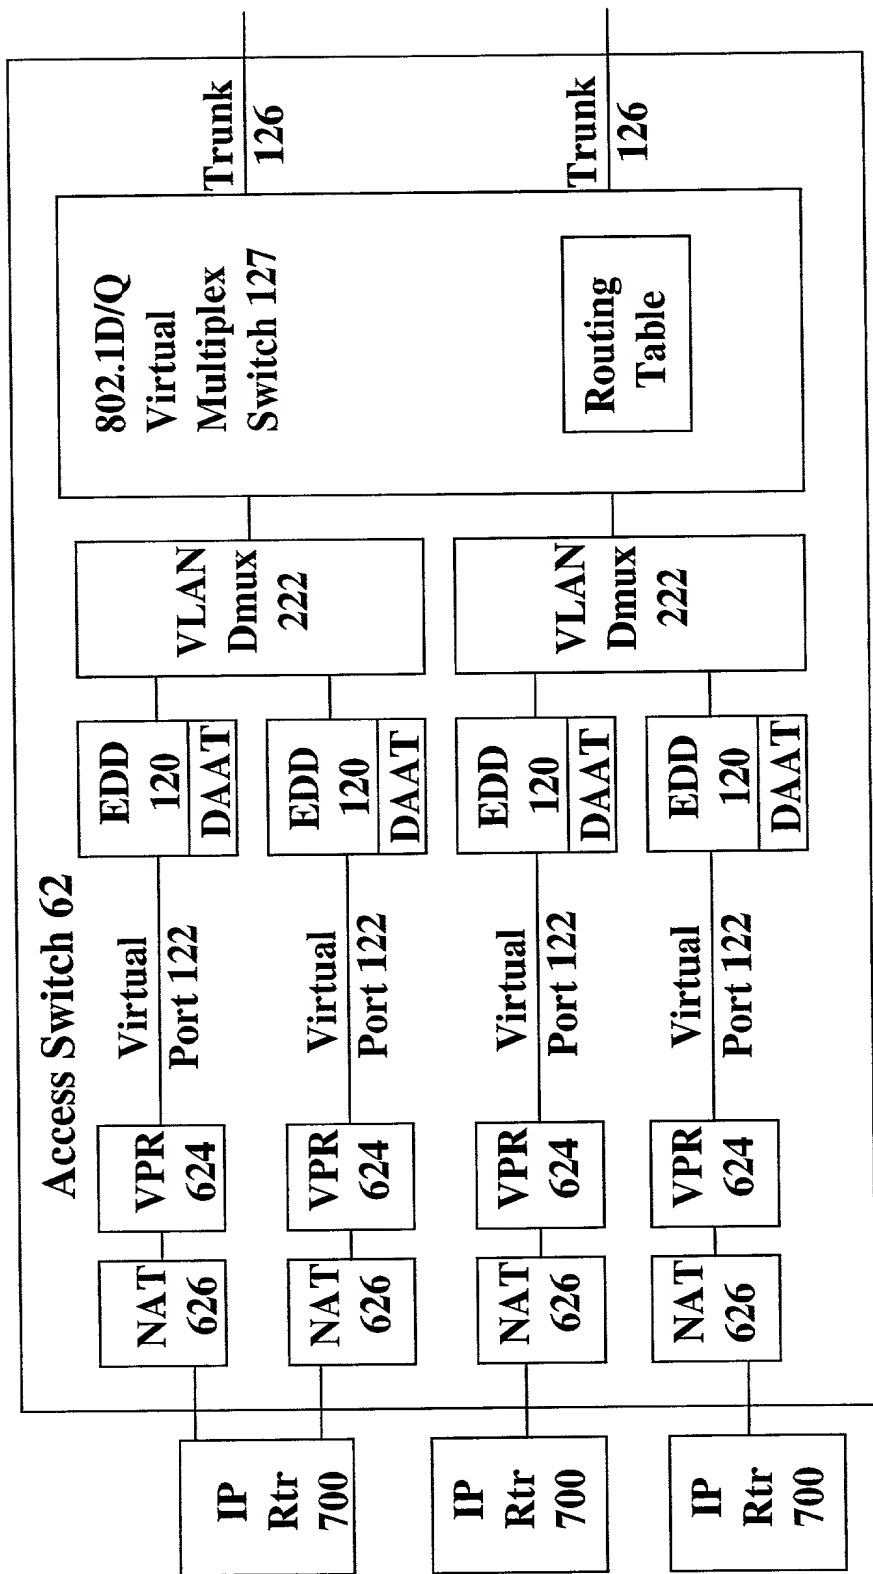


Fig. 9



**Fig. 10**

# DECLARATION FOR PATENT APPLICATION AND APPOINTMENT OF AGENT

Case: 10346RO

As a below-named Inventor, I hereby declare that:

My Residence, Post Office address and Citizenship are as stated below next to my name.

☐ I believe that I am the original, first and sole inventor

☒ I believe I am an original, first and joint inventor

of the subject matter which is claimed and for which a patent is sought on the invention entitled:

## VIRTUAL PRIVATE NETWORKS AND METHODS FOR THEIR OPERATION

the Specification of which

☒ is attached hereto

☐ was filed on \_\_\_\_\_ as U.S. Application or PCT International Application No. \_\_\_\_\_

☐ and was amended on \_\_\_\_\_ (if applicable)

I hereby state that I have reviewed and understand the contents of the above-identified Specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the Examination of the Application in accordance with Title 37, Code of Federal Regulations, §1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, §119(a)-(d) of any foreign Application(s) for Patent or Inventor's Certificate listed below and have also identified below any foreign Application for Patent or Inventor's Certificate having a filing date before that of the Application on which priority is claimed:

### PRIOR FOREIGN APPLICATION(S)

Priority  
Claimed

Number: _____	Country: _____	Date Filed: _____	_____
Number: _____	Country: _____	Date Filed: _____	_____

I hereby claim the benefit under Title 35, United States Code, §119(e) of any United States provisional Application(s) listed below.

Application Number: _____	Date Filed: _____
Application Number: _____	Date Filed: _____

I hereby claim the benefit under Title 35, United States Code, §120 of any United States Application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States Application in the manner provided by the first paragraph of Title 35, United States Code, §112, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations, §1.56 which occurred between the filing date of the prior application and the National or PCT International filing date of the Application.

Application Number: _____	Date Filed: _____	Status: _____
Application Number: _____	Date Filed: _____	Status: _____
Application Number: _____	Date Filed: _____	Status: _____

I hereby appoint **C.W. Junkin** c/o Northern Telecom Limited, Patent Department, P.O. Box 3511 Station C, Ottawa, Ontario, Canada, K1Y 4H7, Registration No. **32,812** and telephone no. (613) 721- 3013 as my Agent to prosecute this application and to transact all business in the Patent and Trademark Office connected therewith.

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true, and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the Application or any Patent issued thereon.

Full Name of First or Sole Inventor: <b>David MacDonald DELANEY</b>	Signature of First Inventor:	Date:
Residence Address: <b>142 Waverley Street, Apartment 2A, Ottawa, Ontario K2P 0V4 CANADA</b>	Country of Citizenship: <b>CANADA</b>	
Post Office Address: <b>same as above</b>		
Full Name of Second Inventor (if any): <b>Peter Martin Kenneth COTTREAU</b>	Signature of Second Inventor:	Date:
Residence Address: <b>5 Links Drive South, R.R. #4, Ashton, Ontario K0A 1B0 CANADA</b>	Country of Citizenship: <b>CANADA</b>	
Post Office Address: <b>same as above</b>		
Full Name of Third Inventor (if any): <b>Alan James HURREN</b>	Signature of Third Inventor:	Date:
Residence Address: <b>23 Antler Avenue, Nepean, Ontario K2J 1Z4 CANADA</b>	Country of Citizenship: <b>CANADA</b>	
Post Office Address: <b>same as above</b>		

Signatures should conform to names as typewritten.

☐ Additional inventors on attached Page 2

Form NTP (06/98)